

Concept Conductor: Orchestrating Multiple Personalized Concepts in Text-to-Image Synthesis

Zebin Yao, Fangxiang Feng*, Ruifan Li, Xiaojie Wang

Beijing University of Posts and Telecommunications
{zebin.yao, fxfeng, rfli, xjwang}@bupt.edu.cn

Abstract

The customization of text-to-image models has seen significant advancements, yet generating multiple personalized concepts remains a challenging task. Current methods struggle with attribute leakage and layout confusion when handling multiple concepts, leading to reduced concept fidelity and semantic consistency. In this work, we introduce a novel training-free framework, Concept Conductor, designed to ensure visual fidelity and correct layout in multi-concept customization. Concept Conductor isolates the sampling processes of multiple custom models to prevent attribute leakage between different concepts and corrects erroneous layouts through self-attention-based spatial guidance. Additionally, we present a concept injection technique that employs shape-aware masks to specify the generation area for each concept. This technique injects the structure and appearance of personalized concepts through feature fusion in the attention layers, ensuring harmony in the final image. Extensive qualitative and quantitative experiments demonstrate that Concept Conductor can consistently generate composite images with accurate layouts while preserving the visual details of each concept. Compared to existing baselines, Concept Conductor shows significant performance improvements. Our method supports the combination of any number of concepts and maintains high fidelity even when dealing with visually similar concepts. The code and trained models will be made publicly available.

Introduction

Text-to-image diffusion models (Nichol et al. 2021; Saharia et al. 2022; Ramesh et al. 2022; Rombach et al. 2022; Podell et al. 2023) have achieved remarkable success in generating realistic high-resolution images. Building on this foundation, techniques for personalizing these models have also advanced. Various methods for single-concept customization (Dong, Wei, and Lin 2022; Ruiz et al. 2023; Gal et al. 2022; Voynov et al. 2023; Alaluf et al. 2023) have been proposed, enabling the generation of images of the target concept in specified contexts based on user-provided visual conditions. These methods allow users to place real-world subjects into imagined scenes, greatly enriching the application scenarios of image generation.

*Corresponding author

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

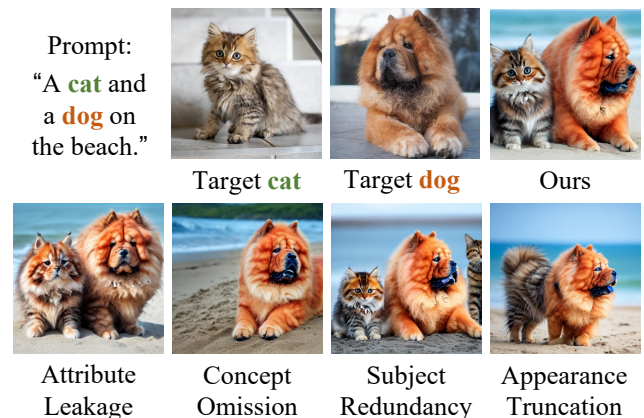


Figure 1: Results from existing multi-concept customization methods (second row) and our method (top right). Our method aims to address attribute leakage and layout confusion (concept omission, subject redundancy, appearance truncation), producing visually faithful and text-aligned images.

Despite the excellent performance of existing methods for single-concept customization, handling multiple concepts remains challenging. Current methods (Kumari et al. 2023; Liu et al. 2023b; Gu et al. 2024) often mix the attributes of multiple concepts or fail to align well with the given text prompts, especially when the target concepts are visually similar (e.g., a cat and a dog). We categorize these failures as attribute leakage and layout confusion. Layout confusion can be further divided into concept omission, subject redundancy, and appearance truncation, as shown in Figure 1. **Attribute leakage** denotes the application of one concept’s attributes to another (e.g., a cat acquiring the fur and eyes of a dog). **Concept omission** indicates one or more target concepts not appearing in the image (e.g., the absence of the target cat). **Subject redundancy** refers to the appearance of extra subjects similar to the target concept (e.g., an extra cat). **Appearance truncation** signifies the target concept’s appearance being observed only in a partial area of the subject (e.g., the upper half of a dog and the lower half of a cat).

To address these challenges, we introduce Concept Con-

ductor, a novel inference framework for multi-concept customization. This framework aims to seamlessly integrate multiple personalized concepts with accurate attributes and layout into a single image based on the given text prompts, as illustrated in Figure 2. Our method comprises three key components: multipath sampling, layout alignment, and concept injection. **Multipath sampling** allows the base model and different single-concept models to retain their independent denoising processes, thereby preventing attribute leakage between concepts. Instead of training a model containing multiple concepts to directly generate the final image, we let each single-concept model first focus on generating its corresponding concept. These generated subjects are then integrated into a single image, avoiding interference and conflict between concepts. **Layout alignment** ensures each model produces the correct layout, fundamentally addressing layout confusion. Specifically, we borrow the layout from a normal image and align the intermediate representations produced by each model in the self-attention layers with it. This reference image is flexible and easy to obtain. It can be a real photo, generated by advanced text-to-image models, or even a simple collage created by the user. **Concept injection** enables each concept to fully inject its visual features into the final generated image, ensuring harmony. We use shape-aware masks to define the generation area for each concept and inject the visual details (including structure and appearance) of personalized concepts through feature fusion in the attention layers. At each step of the denoising process, we first use layout alignment to correct the input latent space representation, and then use concept injection to obtain the next representation. Multipath sampling is implemented in both layout alignment and concept injection to ensure the independence of each subject and coordination between different subjects.

To evaluate the effectiveness of the proposed method, we create a new dataset containing 30 concepts, covering representative categories such as humans, animals, objects, and buildings. We also introduce a fine-grained metric specifically designed for multi-concept customization to measure the visual consistency between the generated subjects and the given concepts. Extensive experiments demonstrate that our method can consistently generate composite images with correct layouts while fully preserving the attributes of each personalized concept, regardless of the number or similarity of the target concepts. Both qualitative and quantitative comparisons highlight the advantages of our method in terms of concept fidelity and alignment with textual semantics.

Our contributions can be summarized as follows:

- We introduce Concept Conductor, a novel framework for multi-concept customization, preventing attribute leakage and layout confusion through multipath sampling and self-attention-based spatial guidance.
- We develop a concept injection technique, utilizing shape-aware masks and feature fusion to ensure harmony and visual fidelity in multi-concept image generation.
- We construct a new dataset containing 30 personalized concepts across representative categories. Comprehensive experiments on this dataset demonstrate the superior

concept fidelity and alignment with textual semantics of the proposed Concept Conductor, including for concepts that are visually similar.

Related Work

Text-to-Image Diffusion Models

In recent years, text-to-image diffusion models have excelled in generating realistic and diverse images, becoming the mainstream approach in this field. Trained on large-scale datasets like LAION(Schuhmann et al. 2022), models such as GLIDE(Nichol et al. 2021), DALL-E 2(Ramesh et al. 2022), Imagen(Saharia et al. 2022), and Stable Diffusion(Rombach et al. 2022) can produce high-quality and text-aligned outputs. However, these models struggle to understand the relationships between multiple concepts, resulting in generated content that fails to fully convey the original semantics. This issue is exacerbated when dealing with visually similar concepts. In this work, we apply our method to the publicly available Stable Diffusion(Rombach et al. 2022), which is based on the Latent Diffusion Model (LDM) architecture. LDM operates in the latent space of a Variational Autoencoder (VAE), iteratively denoising to recover the latent representation of an image from Gaussian noise. At each timestep, the noisy latents z_t are fed into the denoising network ϵ_θ , which predicts the current noise $\epsilon_\theta(z_t, y, t)$ based on the encoded prompt y .

Customization in T2I Diffusion Models

Several works(Ruiz et al. 2023; Gal et al. 2022; Voynov et al. 2023; Alaluf et al. 2023) have customized text-to-image models to generate images of target concepts in new contexts by learning new visual concepts from user-provided example images. For instance, DreamBooth(Ruiz et al. 2023) embeds specific visual concepts into a pre-trained model by fine-tuning its weights. Textual Inversion(Gal et al. 2022) represents new concepts by optimizing a text embedding, later improved by P+(Voynov et al. 2023) and NeTI(Alaluf et al. 2023). Custom Diffusion(Kumari et al. 2023) explores multi-concept customization through joint training, but it requires training a separate model for each combination and often faces severe attribute leakage. Recent works(Liu et al. 2023b; Gu et al. 2024) propose frameworks that combine multiple single-concept models and introduce manually defined layouts in attention maps to aid generation. For example, Cones 2(Liu et al. 2023b) proposes residual embedding-based concept representations for textual composition and emphasizes or de-emphasizes a concept at a specific location by editing cross-attention. Mix-of-show(Gu et al. 2024) merges multiple custom models into one using gradient fusion and restricts each concept’s appearance area through region-controlled sampling. These works cannot fully avoid interference between concepts, leading to low success rates when handling similar concepts or more than two concepts. Additionally, by only manipulating cross-attention and neglecting the impact of self-attention on image structure, these methods often result in mismatched structures and appearances, causing layout control failures. In contrast, our method prevents attribute leakage by isolating the

sampling processes of different single-concept models and achieves stable layout control through self-attention-based spatial guidance.

Spatial Control in T2I Diffusion Models

Using text prompts alone is insufficient for precise control over image layout or structure. Some methods(Avrahami et al. 2023; Li et al. 2023; Zhang, Rao, and Agrawala 2023; Mou et al. 2024) introduce layout conditions by training additional modules to generate controllable images. For example, GLIGEN(Li et al. 2023) adds trainable gated self-attention layers to allow extra input conditions, such as bounding boxes. To achieve finer spatial control, Control-Net(Zhang, Rao, and Agrawala 2023) and T2I-Adapter(Mou et al. 2024) introduce image-based conditions like keypoints, sketches, and depth maps by training U-Net encoder copies or adapters. These methods can stably control image structure but limit the target subjects’ poses, and these conditional images are difficult for users to create. Some training-free methods achieve spatial guidance by manipulating attention layers during sampling. Most works(Ma et al. 2024; Kim et al. 2023; He, Salakhutdinov, and Kolter 2023) attempt to alleviate attribute leakage and control layout by directly editing attention maps but have low success rates. Several gradient-based methods (Couairon et al. 2023; Xie et al. 2023; Phung, Ge, and Huang 2024) calculate the loss between attention and the given layout and introduce layout information into the latent space representation by optimizing the loss gradients. These methods align the generated image layout with coarse visual conditions (e.g., bounding boxes and semantic segmentation maps), often resulting in high-frequency detail loss and image quality degradation, even with complex loss designs. In this work, we propose extracting layout information from an easily obtainable reference image as a supervisory signal, which not only stably controls the layout but also preserves the diversity of the generated subjects’ poses while avoiding image distortion.

Method

Preliminary: Attention Layers in LDM

LDM employs a U-Net as the denoising network, consisting of a series of convolutional layers and transformer blocks. In each block, intermediate features produced by the convolutional layers are passed to a self-attention layer followed by a cross-attention layer. Given an input feature h_{in} , the output feature in each attention layer is computed as $h_{out} = AV$, where $A = softmax(QK^T)$. Here, $Q = f_Q(h_{in})$, $K = f_K(c)$, and $V = f_V(c)$ are obtained through learned projectors f_Q , f_K , and f_V , with $c = h_{in}$ for self-attention and $c = y$ for cross-attention. Self-attention enhances the quality of the generated image by capturing long-range dependencies in the image features, while cross-attention integrates textual information into the generation process, enabling the generated image to reflect the content of the text prompt. Furthermore, extensive researches (Liu et al. 2024; Patashnik et al. 2023; Hertz et al. 2022) have shown that self-attention controls the structure of the image (e.g., shapes, geometric relationships), whereas cross-attention controls the




Figure 2: Overview of our proposed Concept Conductor. At each denoising step, the input latent vector z_t is first corrected to z_t' by the Layout Alignment module. z_t' is then sent to the Concept Injection module for denoising, producing the next latent vector z_{t-1} . Both Layout Alignment and Concept Injection utilize the Multipath Sampling structure. After denoising, our method can generate images that align with the given text prompt and visual concepts.

appearance of the image (e.g., colors, materials, textures).

Preliminary: ED-LoRA

ED-LoRA is a method for single-concept customization, primarily involving learnable hierarchical text embeddings and low-rank adaptation (LoRA) applied to pre-trained weights. To learn the representation of a concept within the pre-trained model’s domain, it creates layer-wise embeddings for the target concept’s token following P+(Voynov et al. 2023). Additionally, to capture out-of-domain visual details, it fine-tunes the pre-trained text encoder and U-Net using LoRA(Hu et al. 2021). In this paper, ED-LoRA is used as our single-concept model by default.

Overview of Concept Conductor

Our method comprises three components: multipath sampling, layout alignment, and concept injection, as illustrated in Figure 2. At each denoising step, we first correct the input latents z_t through layout alignment, obtaining new latents z_t' that carry the layout information from the reference image. Then, we inject the personalized concepts from the custom models into the base model and denoise z_t' to obtain the next latents z_{t-1} . Multipath sampling is implemented in both layout alignment and concept injection to ensure the independence of each concept and coordination between different subjects.

Multipath Sampling

Joint training or model fusion methods often lead to attribute leakage between different concepts (as shown in Figure 1) and require additional optimization steps for each combination. To directly utilize multiple existing single-concept models for composite generation without attribute leakage, we propose a multipath sampling structure. This structure incorporates a base model ϵ_{θ}^{base} and multiple custom models $\epsilon_{\theta}^{V_i}$ (implemented with ED-LoRA(Gu et al. 2024)) for personalized concepts V_i , as illustrated in Figure 3.

Given several custom models $\epsilon_{\theta}^{V_i}$ and a text prompt p , at each timestep t , we maintain the independent denoising pro-

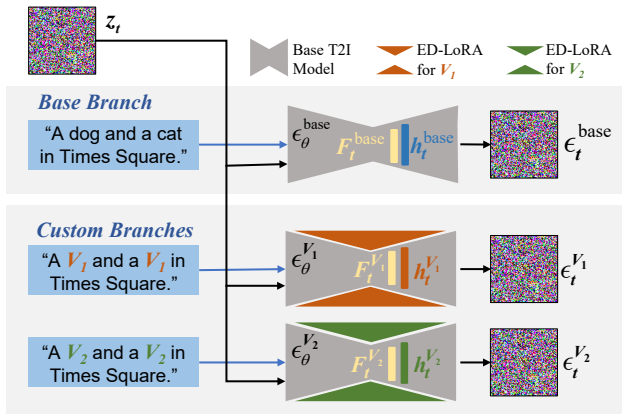


Figure 3: Illustration of multipath sampling. custom models $\epsilon_{\theta}^{V_1}$ and $\epsilon_{\theta}^{V_2}$ are created by adding ED-LoRA to the base model ϵ_{θ}^{base} . The base prompt and edited prompts are sent to the base model and custom models, respectively. Different models receive the same latent input z_t and predict different noises. Self-attention features F_t^{base} , $F_t^{V_1}$, $F_t^{V_2}$, and the output feature maps of the attention layers h_t^{base} , $h_t^{V_1}$, $h_t^{V_2}$ are recorded during this process.

cess for each model: $\epsilon_t^{V_i} = \epsilon_{\theta}^{V_i}(z_t, t, p)$. When the prompt contains similar concepts, models may struggle to distinguish them, leading to attribute leakage. Therefore, we edit the input text prompt for each custom model to help them focus on generating the corresponding single concept. Given a base prompt p_{base} , we replace tokens visually similar to the target concept V_i with tokens representing the target concept, creating a prompt variant p_{V_i} . For example, for the prompt “A dog and a cat on the beach” and concepts of a dog V_1 and a cat V_2 , we edit the text to obtain two modified prompts: p_{V_1} = “A $\langle V_1 \rangle$ and a $\langle V_1 \rangle$ on the beach” and p_{V_2} = “A $\langle V_2 \rangle$ and a $\langle V_2 \rangle$ on the beach”.

After editing the prompts, the denoising process for the custom models can be expressed as: $\epsilon_t^{V_i} = \epsilon_{\theta}^{V_i}(z_t, t, p_{V_i})$. Meanwhile, the base prompt is sent to the base model to retain global semantics: $\epsilon_t^{base} = \epsilon_{\theta}^{base}(z_t, t, p_{base})$. Through multipath sampling, each custom model receives only information relevant to its corresponding concept, fundamentally preventing attribute leakage between different concepts.

Layout Alignment

Existing multi-concept customization methods often suffer from layout confusion (as shown in Figure 1), especially when the target concepts are visually similar or numerous. To address these challenges, we introduce a reference image to correct the layout during the generation process. For example, to generate an image of a specific dog and cat in a specific context, we only need a reference image containing an ordinary cat and dog. One simple approach to achieve layout control is to convert the reference image into abstract visual conditions (e.g., keypoints or sketches) and then use ControlNet(Zhang, Rao, and Agrawala 2023) or T2I-Adapter(Mou et al. 2024) for spatial guidance, which

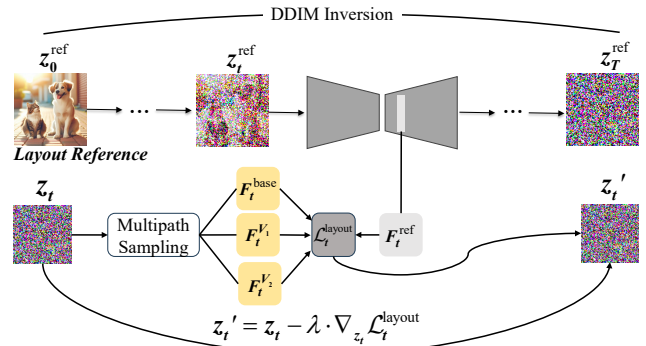


Figure 4: Illustration of layout alignment. The self-attention feature F_t^{ref} of the layout reference image is extracted through DDIM inversion, which is then used to compute the loss with F_t^{base} , $F_t^{V_1}$, and $F_t^{V_2}$, updating the input latent vector z_t . For simplicity, the conversion from pixel space to latent space is omitted.

limits variability and flexibility and may reduce the fidelity of the target concepts. Another approach is to directly inject the full self-attention of the reference image into the generation process, transferring the overall structure of the image(Hertz et al. 2022). This strictly limits the poses of the target subjects, reducing the diversity of the generated images. Additionally, it requires structural similarity between the reference image subjects and target concepts to avoid distortions caused by shape mismatches.

To align the layout while preserving the structure of the target concepts, we propose a gradient-guided approach, as shown in Figure 4. Given a reference image, we perform DDIM inversion to obtain the latent space representation z_t^{ref} and record the self-attention features F_t^{ref} at each timestep. Similarly, we record the self-attention features of the base model and each custom model during the generation process, denoted as F_t^{base} and $F_t^{V_i}$, respectively, as shown in Figure 2. An optimization objective is set to encourage the generated image’s layout to align with the given layout:

$$\mathcal{L}_t^{layout} = \|F_t^{base} - F_t^{ref}\|_2 + \alpha \frac{1}{N} \sum_{i=1}^N \|F_t^{V_i} - F_t^{ref}\|_2 \quad (1)$$

where α represents the weighting coefficient, and N denotes the number of personalized concepts. We use gradient descent to optimize this objective, obtaining the corrected latent space representation:

$$z_t' = z_t - \lambda \cdot \nabla_{z_t} \mathcal{L}_t^{layout} \quad (2)$$

where λ represents the gradient descent step size. Through layout alignment, we ensure that the generated image mimics the reference image’s layout without any confusion.

Concept Injection

After layout alignment, the original latents z_t are replaced with the corrected latents z_t' , and multipath sampling is used to generate the next latents z_{t-1} . The goal is to inject

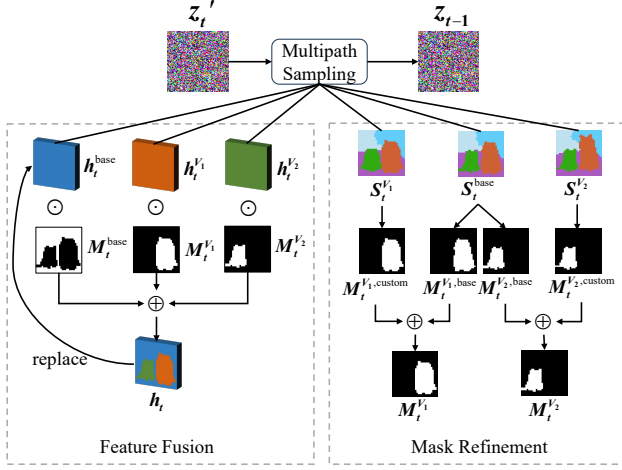


Figure 5: Illustration of concept injection, consisting of two parts: (1) Feature Fusion. The output feature maps of the attention layers from different models are multiplied by their corresponding masks and summed to obtain the fused feature map h_t , which is used to replace the original feature map h_t^{base} . (2) Mask Refinement. Segmentation maps are obtained by clustering on the self-attention, and the masks required for feature fusion are extracted from these maps.

the subjects generated by different custom branches into the base branch to create a composite image. A naive way is to spatially fuse the noise predicted by different models:

$$\epsilon_t^{fuse} = \epsilon_t^{base} \odot M^{base} + \sum_{i=1}^N \epsilon_t^{V_i} \odot M^{V_i} \quad (3)$$

where ϵ_t^{base} represents the noise predicted by the base model, $\epsilon_t^{V_i}$ represents the noise predicted by the custom model for concept V_i , and M^{base} and M^{V_i} are predefined masks. This method ensures the fidelity of the target concepts but often results in disharmonious images.

To address this issue, we propose an attention-based concept injection technique, including feature fusion and mask refinement, as shown in Figure 5. Spatial fusion is implemented on the output feature maps of all attention layers in the U-Net decoder, as self-attention controls the structure of the subjects and cross-attention controls their appearance, both crucial for reproducing the attributes of the target concepts. For each selected attention layer, the fused output feature is computed as:

$$h_t = h_t^{base} \odot M_t^{base} + \sum_{i=1}^N h_t^{V_i} \odot M_t^{V_i} \quad (4)$$

where $M_t^{base} = 1 - \bigcup_{i=1}^N M_t^{V_i}$. Here, h_t^{base} and $h_t^{V_i}$ represent the output features of the attention layers of the base model and the custom models, respectively, and $M_t^{V_i}$ represents the binary mask of concept V_i at timestep t , specifying the dense generation area of the target concept. The fused feature h_t

is then sent back to the corresponding position in the base model to replace h_t^{base} and complete the denoising process.

Since the poses of the generated subjects are uncertain, predefined masks may not precisely match the shapes and contours of the target subjects, leading to incomplete appearances. To address this, we use mask refinement to allow the masks to adjust according to the shapes and poses of the target subjects during the generation process. Inspired by local-prompt-mixing (Patashnik et al. 2023), we use self-attention-based semantic segmentation to obtain the masks of the target subjects. For each target concept V_i , we cluster the self-attention $A_t^{V_i}$ of the custom model ϵ_{V_i} to obtain a semantic segmentation map $S_t^{V_i}$ and extract the subject’s mask $M_t^{V_i, custom}$. We perform the same operation on the self-attention A_t^{base} of the base model ϵ_{base} , obtaining the semantic segmentation map S_t^{base} and several masks $M_t^{V_i, base}$, $i \in [1, N]$, each corresponding to a subject V_i . To reconcile the shape differences between the subjects in the base model and the custom models, the corresponding masks are merged: $M_t^{V_i} = M_t^{V_i, custom} \cup M_t^{V_i, base}$.

For initialization, we perform DDIM inversion on the reference image and extract the original masks $M_T^{V_i}$ from the self-attention layers in the same way. An alternative way is to use grounding models (e.g., Grounding DINO (Liu et al. 2023a)) and segmentation models (e.g., SAM (Kirillov et al. 2023)) to extract masks in the pixel space, which provides higher resolution masks but requires additional computation and storage overhead. Through concept injection, we ensure the harmony of the image while fully preserving the attributes of the target concepts.

Experiments

Dataset

We construct a dataset covering representative categories such as humans, animals, objects, and buildings, including 30 personalized concepts. Real and anime human images are collected from the Mix-of-show dataset (Gu et al. 2024), while other categories are sourced from the DreamBooth dataset (Ruiz et al. 2023) and CustomConcept101 (Kumari et al. 2023), with 3-15 images per concept. For quantitative evaluation, we select 10 pairs of visually similar concepts and generate 5 text prompts for each pair using ChatGPT (OpenAI 2023). We produce 8 samples for each text prompt using the same set of random seeds, resulting in a total of $10 \times 5 \times 8 = 400$ images per method.

Implementation Details

Our method is implemented on Stable Diffusion v1.5, using images generated by SDXL (Podell et al. 2023) as layout references. In layout alignment, the key of the first self-attention layer in the U-Net decoder is used as the layout feature F_t . For mask refinement, we cluster the attention probabilities in the sixth self-attention layer of the U-Net decoder to extract semantic segmentation maps and scale them to different sizes for feature fusion in different attention layers. In all experiments, the weighting coefficient α is set to 1, and the gradient descent step size λ is set to 10.




Figure 6: Qualitative comparison of multi-concept customization methods. The results show that our method ensures visual fidelity and correct layout, while other methods suffer from severe attribute confusion and layout disorder.

Baselines

We compare our method with three multi-concept customization methods: Custom Diffusion, Cones 2, and Mix-of-Show. For Custom Diffusion, we use the diffusers(von Platen et al. 2022) version implementation, while for the other methods, we use their official code implementations. All experimental settings follow the official recommendations. For Cones 2 and Mix-of-Show, grounding models are used to extract bounding boxes of target subjects from the layout reference image to ensure consistent spatial conditions. To ensure a fair comparison, no additional control models like ControlNet or T2I-Adapter are used.

Evaluation Metrics

We evaluate multi-concept customization methods from two perspectives: text alignment and image alignment. For text alignment, we report results on CLIP(Radford et al. 2021) and ImageReward(Xu et al. 2024). For image alignment, we introduce a new metric called Segmentation Similarity (SegSim) to address the limitations of traditional image similarity methods, which cannot reflect attribute leakage and layout conflicts. SegSim evaluates fine-grained fidelity by using text-guided grounding models and segmentation models to extract subject segments from generated and reference images, then calculating their similarity. We use CLIP(Radford et al. 2021) and DINOv2(Oquab et al. 2023) to calculate segment similarity and report image alignment based on these models. To systematically evaluate omission and redundancy in multi-concept generation, grounding models are used to automatically count the number of target category subjects in each generated image.



Figure 7: Qualitative comparison in challenging scenarios. Mix-of-Show struggles to handle more than two similar concepts or complex layouts, whereas our method demonstrates robust performance even in these complex scenarios.

Qualitative Comparison

We evaluate our method and all baselines on various combinations of similar concepts, as shown in Figure 6. Custom Diffusion and Cones 2 struggle to retain the visual details of target concepts (e.g., the cat’s fur pattern and the backpack’s design), exhibiting severe attribute leakage (e.g., two identical boots) and layout confusion (e.g., missing or redundant teddy bears). Mix-of-Show demonstrates higher fidelity and

Method	Text-alignment		Image-alignment		Counting	
	CLIP-T	IR	CLIP-I	DINO	$n < 2$	$n > 2$
CD	0.2939	0.0133	0.8333	0.6829	0.2750	0.1050
Cones2	0.2954	0.0973	0.8415	0.6928	0.3600	0.0650
MoS	0.2962	0.3944	0.8928	0.7961	0.3400	0.0400
Ours	0.3107	1.2542	0.9190	0.8569	0.0075	0.0350

Table 1: Quantitative Comparison of Multi-Concept Customization Methods. IR stands for ImageReward, CD for Custom Diffusion, and MoS for Mix-of-Show. $n < 2$ indicates omission, while $n > 2$ indicates redundancy.

mitigates attribute leakage but still faces significant concept omission (e.g., missing cartoon backpack) and appearance truncation (e.g., stitched teddy bear). In contrast, our Concept Conductor generates all target concepts with high fidelity without leakage through multipath sampling and concept injection, ensuring correct layout through layout alignment. Our method maintains stable performance across different concept combinations, even when the target concepts are very similar, such as two teddy bears.

We further explore more challenging scenarios and compare our method with Mix-of-Show, as shown in Figure 7. When handling three similar concepts, Mix-of-Show exhibits severe attribute leakage (e.g., both men wearing glasses) and concept omission (e.g., one person missing). Additionally, Mix-of-Show struggles with dense layouts, often resulting in appearance truncation when faced with complex spatial relationships (e.g., the upper half of a cat and the lower half of a dog stitched together). In contrast, our method maintains the visual features of each concept without attribute leakage and faithfully reflects the layout described in the text, even in these complex scenarios.

Quantitative Comparison

As reported in Table 1, our Concept Conductor significantly outperforms previous methods in both image alignment and text alignment. The improvement in image alignment indicates that our method can preserve the visual details of each concept without attribute leakage, primarily due to our proposed multipath sampling framework and attention-based concept injection. The improvement in text alignment is mainly because our method effectively avoids the layout confusion that leads to unfaithful or disharmonious images through layout alignment, thereby enhancing text-image consistency. The significant reduction in omission and redundancy rates also supports this.

Ablation Study

To verify the effectiveness of the proposed components, we conduct qualitative comparisons of various settings, as shown in Figure 8. In Figure 8(a), removing layout alignment leads to incorrect layouts, including appearance truncation (e.g., two dogs incorrectly stitched together) and concept omission (e.g., missing turquoise cup). Figures 8(b) and 8(c) show that disabling either self-attention or cross-attention features during concept injection results in a decline in fidelity, indicating that both are crucial for preserv-




Figure 8: Qualitative comparison of ablation variants. (a) Results without layout alignment (LA). (b) Results without self-attention (SA) features in concept injection. (c) Results without cross-attention (CA) features in concept injection. (d) Results without mask refinement (MR) in concept injection. (e) Results of our complete model.

Method	Text-alignment		Image-alignment		Counting	
	CLIP-T	IR	CLIP-I	DINO	$n < 2$	$n > 2$
w/o LA	0.3020	0.8926	0.9021	0.8258	0.0575	0.0625
w/o SA	0.3095	1.2224	0.8940	0.7820	0.0450	0.0475
w/o CA	0.3027	1.2059	0.9051	0.8194	0.0275	0.0550
w/o MR	0.3034	1.2140	0.9148	0.8478	0.0100	0.0375
Ours	0.3107	1.2542	0.9190	0.8569	0.0075	0.0350

Table 2: Quantitative Comparison of Ablation Variants. Ablation is performed on four components: Layout Alignment (LA), Self-Attention (SA), Cross-Attention (CA), and Mask Refinement (MR).

ing the visual details of the target concepts. Figure 8(d) demonstrates that ablating mask refinement can lead to a mismatch between the generated subject contours and the target concepts, resulting in incomplete appearances (e.g., chow chow’s fur, corgi’s ears, green cup’s handle). To avoid randomness, we conduct quantitative comparisons of various settings using the same data and evaluation metrics as in the previous section, with results reported in Table 2. As shown in Table 2, layout alignment effectively avoids omission and redundancy, significantly improving the alignment of generated images with textual semantics. Feature fusion in both self-attention and cross-attention layers leads to higher image alignment, as both are crucial for reproducing the attributes of the target concepts. Mask refinement further improves text alignment and image alignment by optimizing the edge details of the generated subjects.

Conclusion

We introduce Concept Conductor, a novel inference framework designed to generate realistic images containing multiple personalized concepts. By employing multipath sampling and layout alignment, we addressed the common issues of attribute leakage and layout confusion in multi-concept personalization. Additionally, concept injection is used to create harmonious composite images. Experimental results demonstrate that Concept Conductor can consistently generate composite images with correct layouts, fully preserving the attributes of each concept, even when the target concepts are highly similar or numerous.

Acknowledgments

This work was supported by the NSFC (No.62076032), Beijing Natural Science Foundation (L247010), NSFC (No.61906018) and the CCF-Zhipu Large Model Innovation Fund (NO. CCF-Zhipu202407).

References

- Alaluf, Y.; Richardson, E.; Metzer, G.; and Cohen-Or, D. 2023. A neural space-time representation for text-to-image personalization. *ACM Transactions on Graphics (TOG)*, 42(6): 1–10.
- Avrahami, O.; Hayes, T.; Gafni, O.; Gupta, S.; Taigman, Y.; Parikh, D.; Lischinski, D.; Fried, O.; and Yin, X. 2023. Spatext: Spatio-textual representation for controllable image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18370–18380.
- Couairon, G.; Careil, M.; Cord, M.; Lathuiliere, S.; and Verbeek, J. 2023. Zero-shot spatial layout conditioning for text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2174–2183.
- Dong, Z.; Wei, P.; and Lin, L. 2022. Dreamartist: Towards controllable one-shot text-to-image generation via positive-negative prompt-tuning. *arXiv preprint arXiv:2211.11337*.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*.
- Gu, Y.; Wang, X.; Wu, J. Z.; Shi, Y.; Chen, Y.; Fan, Z.; Xiao, W.; Zhao, R.; Chang, S.; Wu, W.; et al. 2024. Mix-of-show: Decentralized low-rank adaptation for multi-concept customization of diffusion models. *Advances in Neural Information Processing Systems*, 36.
- He, Y.; Salakhutdinov, R.; and Kolter, J. Z. 2023. Localized text-to-image generation for free via cross attention control. *arXiv preprint arXiv:2306.14636*.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Kim, Y.; Lee, J.; Kim, J.-H.; Ha, J.-W.; and Zhu, J.-Y. 2023. Dense text-to-image generation with attention modulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7701–7711.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1931–1941.
- Li, Y.; Liu, H.; Wu, Q.; Mu, F.; Yang, J.; Gao, J.; Li, C.; and Lee, Y. J. 2023. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22511–22521.
- Liu, B.; Wang, C.; Cao, T.; Jia, K.; and Huang, J. 2024. Towards Understanding Cross and Self-Attention in Stable Diffusion for Text-Guided Image Editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7817–7826.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023a. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Liu, Z.; Zhang, Y.; Shen, Y.; Zheng, K.; Zhu, K.; Feng, R.; Liu, Y.; Zhao, D.; Zhou, J.; and Cao, Y. 2023b. Cones 2: Customizable image synthesis with multiple subjects. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 57500–57519.
- Ma, W.-D. K.; Lahiri, A.; Lewis, J. P.; Leung, T.; and Kleijn, W. B. 2024. Directed diffusion: Direct control of object placement through attention guidance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4098–4106.
- Mou, C.; Wang, X.; Xie, L.; Wu, Y.; Zhang, J.; Qi, Z.; and Shan, Y. 2024. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4296–4304.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- OpenAI. 2023. ChatGPT. <https://openai.com/chatgpt/>.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Patashnik, O.; Garibi, D.; Azuri, I.; Averbuch-Elor, H.; and Cohen-Or, D. 2023. Localizing object-level shape variations with text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23051–23061.
- Phung, Q.; Ge, S.; and Huang, J.-B. 2024. Grounded text-to-image synthesis with attention refocusing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7932–7942.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22500–22510.

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. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494.

Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35: 25278–25294.

von Platen, P.; Patil, S.; Lozhkov, A.; Cuenca, P.; Lambert, N.; Rasul, K.; Davaadorj, M.; Nair, D.; Paul, S.; Berman, W.; Xu, Y.; Liu, S.; and Wolf, T. 2022. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>.

Voyinov, A.; Chu, Q.; Cohen-Or, D.; and Aberman, K. 2023. p+: Extended textual conditioning in text-to-image generation. *arXiv preprint arXiv:2303.09522*.

Xie, J.; Li, Y.; Huang, Y.; Liu, H.; Zhang, W.; Zheng, Y.; and Shou, M. Z. 2023. Boxdiff: Text-to-image synthesis with training-free box-constrained diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7452–7461.

Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2024. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.