

Concept Weaver: Enabling Multi-Concept Fusion in Text-to-Image Models

Gihyun Kwon¹ Simon Jenni² Dingzeyu Li² Joon-Young Lee²
 Jong Chul Ye¹ Fabian Caba Heilbron²

KAIST¹ Adobe²

[gihyun, jong.ye]@kaist.ac.kr [jenni, dinli, jolee, caba]@adobe.com

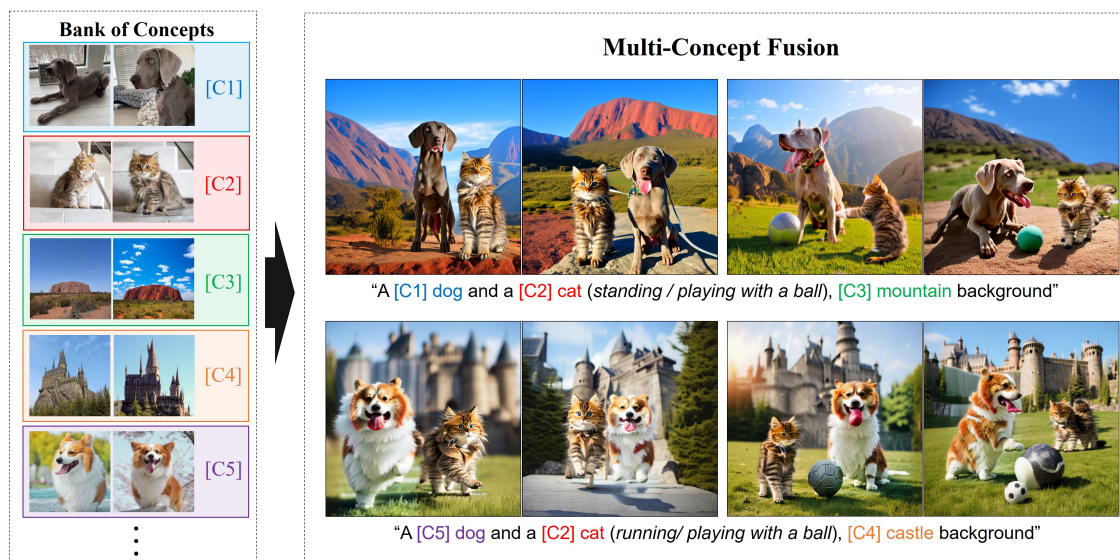


Figure 1. **Concept Weaver’s Generation Results.** Our method, Concept Weaver, can inject the appearance of arbitrary off-the-shelf concepts (from a Bank of Concepts) to generate realistic images.

Abstract

While there has been significant progress in customizing text-to-image generation models, generating images that combine multiple personalized concepts remains challenging. In this work, we introduce *Concept Weaver*, a method for composing customized text-to-image diffusion models at inference time. Specifically, the method breaks the process into two steps: creating a template image aligned with the semantics of input prompts, and then personalizing the template using a concept fusion strategy. The fusion strategy incorporates the appearance of the target concepts into the template image while retaining its structural details. The results indicate that our method can generate multiple custom concepts with higher identity fidelity compared to alternative approaches. Furthermore, the method is shown to seamlessly handle more than two concepts and

closely follow the semantic meaning of the input prompt without blending appearances across different subjects.

1. Introduction

Text-to-image generation models have shown impressive capabilities [21, 23, 28] in the last few years. Existing open source [21] and commercial solutions such as Adobe Firefly have enabled aspiring creatives to generate images with unprecedented quality by simply crafting text prompts. Progress has also been attained in developing models that can customize images for your own subjects or visual concepts [3, 11, 22, 25]. These technologies have opened the door for new ways of content creation, where aspiring creators can craft stories with personalized characters under different scenes and styles.

While there has been significant progress in customizing text-to-image generation models, generating images that

This work is done when Gihyun Kwon was an intern at Adobe Research.

combine multiple personalized concepts remains challenging. Several approaches [11, 25] offer the ability to jointly train models for multiple concepts or merge customized models, enabling the creation of scenes with more than one personalized concept. However, it often fails to generate semantically related concepts (*e.g.*, cat and dog) and struggles to scale beyond three or more concepts. More recently, Mix-of-show [4] has addressed the issue of multi-concept generation with disentangled Low-Rank (LoRa) [9] weight merging and regional guidance at the sampling stage. However, the model still suffers from mixed concepts due to the difficulty of weight merging.

In this paper, we propose a tuning-free method for composing customized text-to-image diffusion models at inference time. We illustrate our key idea in Figure 2, where the goal is to generate images featuring more than two custom concepts. Specifically, rather than generating a personalized image from scratch, we break the process into two steps: first, we create a template image that aligns with the semantics of the input prompt, and then we personalize this template image using a novel concept fusion strategy. The fusion strategy takes as input the non-personalized template image along with region concept guidance (obtained automatically) to generate an edited image that retains the template’s structural details while incorporating the target concepts’ appearance and style. This fusion approach injects concept details into specific spatial regions, allowing us to compose multiple concepts (from the Bank of Concepts) in generated images without blending appearances across different subjects.

Our empirical evaluations show that the proposed method is able to generate multiple custom concepts with higher concept fidelity. In particular, as shown in Section 4, we observe that our method can compose images without blending appearances for semantically related concepts (cats and dogs). Second, we notice that our model can seamlessly handle more than two concepts, *e.g.*, two subjects and a custom background, while the baseline approaches struggle. Finally, we find that the images generated by our method closely follow the semantic meaning of the input prompt achieving high CLIP scores [11]. Ours also has robustness on architecture as it can be used in both of full fine-tuning and Low-Rank adaptation, which is more efficient in computation.

2. Related Work

Text-to-image Diffusion Models. Text-to-image generation models have made significant progress, starting from early GAN-based models [2, 29] to recent diffusion-based models [20, 21, 23, 28]. Various open source models and commercial models like Adobe Firefly have contributed to this development. The recent introduction of Stable Diffusion models [21] has led to the exploration of various appli-

cations such as mask-based image editing [1], image translation [16, 18, 26], and style transfer based on text [30]. Moreover, the attention-based structure of stable diffusion has inspired different editing methods [7, 17, 26].

Diffusion Model Customization. Building on the advancements of these T2I models, research on customizing T2I models using user-prepared images or visual concepts has gained attention. The seminal work of Textual Inversion [3] has focused on finding optimized textual embeddings for custom concepts to generate concept-reflecting images. Subsequent research has improved performance by finding extended textual embeddings [12, 27] or fine-tuning model parameters [11, 22], enabling more efficient and flexible customization.

Extended from the previous single-concept frameworks, customization involving multiple concepts has also been attempted. These approaches include methods using joint training for simultaneously embedding the multi-concepts [5, 11], weight merging of single-concept customized model parameters [11, 25], and spatial guidance [13]. However, these approaches face challenges when the number of concepts increases or when the semantic distance between the concepts is close, resulting in the disappearance or blending of specific concepts. To address this, recent work of Mix-of-show [4] applies regional guidance during the sampling process using merged weights to resolve the issue of concept blending. However, the approach still requires additional optimization steps for weight merging and may experience fluctuations in quality due to the sensitivity to regional guidance.

3. The Concept Weaver’s Method

In this section, we introduce Concept Weaver, an innovative method designed to generate high-quality images that incorporate multiple custom concepts. Traditional models often struggle with generating complex, multi-concept images in a single step. Concept Weaver addresses this by employing a cascading generation process, which we illustrate in Figure 2. Consider the prompt: “A [C1]dog and a [C2]cat playing with a ball, [C3]mountain background”, where [C1, C2, C3] denote custom concepts. Our approach begins by personalizing text-to-image models for each concept (Step 1). Next, we select a non-personalized ‘template image’ using the given prompt, either from a text-to-image model or a real-world source (Step 2). In the third step, we extract latent representations from this template to aid in later editing. The fourth step involves identifying and isolating the specific regions of the template image that correspond to the target subjects. Finally, our key contribution (Step 5) combines these latent representations, targeted spatial regions, and personalized models to reconstruct the




Figure 2. **Concept Weaver’s Method.** First, we fine-tune a text-to-image model for each target concept in the bank (Step 1). Then we source a template image (Step 2). Given the template image, we apply the inversion process with simultaneous feature extraction to save its structural information (Step 3). In Step 4, we extract region masks from the template image with off-the-shelf models [10]. With extracted features and masks, we generate the multi-concept image in Step 5.

template image, infusing it with the specified concepts. We present each of these key steps in detail next.

Step 1: Concept Bank Training. In this step we fine-tune a pretrained text-to-image model to embed each of the target concepts in the bank. Among the various customization strategies, we leverage Custom Diffusion [11] as it does not change any residual network or self-attention layers. In practice, Custom Diffusion only fine-tunes the cross-attention layers of the U-Net model ϵ_θ . Specifically, with the text condition $p \in R^{s \times d}$ and self-attention feature $f \in R^{(h \times w) \times c}$, the cross attention layer consists of $Q = W^q f, K = W^k p, V = W^v p$.

We only fine-tune the ‘key’ and the ‘value’ weight parameters W^k, W^v of the cross-attention layers. Also, we use modifier tokens [V*], which are placed ahead of the concept word (e.g., [V*] dog) and operate as a constraint to general concepts. We augment the fine-tuning process with robust data augmentation techniques. Since we can incorporate an arbitrary personalization approach if the method is only related to cross-attention layers, we can naturally extend the approach to an efficient LoRA [9]-based fine-tuning method. We will show the flexibility of the proposed approach in our experiment part.

Step 2 : Template Image Generation. One of our key insights is to cascade the multi-concept generation process – we start from a template image that can be customized/personalized with the target concepts in the given prompt. To source a template image we can rely on existing text-to-image models but also on real images if given. They should include the semantic objects (or characters) with specific background desired in the prompt. In practice, we generate template images using Stable Diffusion [21] model version ≥ 2.0 .

Step 3 : Inversion and Feature Extraction. After sourcing a template image, we apply an inversion process to obtain a latent representation that will help guide our generation

process. In this stage, we borrow the image inversion and feature extraction schemes proposed in plug-and-play diffusion (PNP) [26]. More specifically, as shown in Figure 3 (a), from the source image x_{src} we generate the noisy latent space z_T with the DDIM [24] forward process. From the inverted latent z_T , we can accurately reconstruct the source image using a reverse DDIM process [24]. We provide more details about the inversion process in the supplementary material. During the reverse reconstruction process, we extract the features from the U-Net’s l -th layer f_t^l at each timestep t . These features include intermediate outputs from residual layers and self-attention activations. As proposed in PNP diffusion, we extract the ResNet output from $l = 4$ and self-attention maps from $l = 4, 7, 9$. Inspired by the recent negative prompt inversion [6], we used the reference text condition p_{src} during the inversion process.

Step 4 : Mask Generation. Given an inverted latent and pre-calculated features, we can guide the structural information of the subsequent generation process. However, we using the structural guidance cannot guarantee the concept-wise editing of each targeting concepts and generated images often yields mixed concepts. Therefore, we use the masked guidance in which we apply the personalized generation model to the specific regions which already contains the template objects. In order to obtain the semantic mask regions, we leveraged the Segment Anything Model [10]. To further avoid the manual seeding of segmentation model, we incorporated the pre-trained text conditional grounding model [15] to obtain the bounding box regions with given text prompts. We then obtain the box regions giving single concept-wise words such as ‘a dog’, ‘a cat’, etc. For N different concepts, we extract concept-wise masks M_1, M_2, \dots, M_N , and set the unmasked region as background mask $M_{bg} = (M_1 \cup M_2 \cup \dots \cup M_N)^c$.

We empirically discovered that when we use directly obtained densely annotated masks, the final output often yields deformed outputs. Therefore instead of using densely annotated mask, we used dilated mask in which the mask region

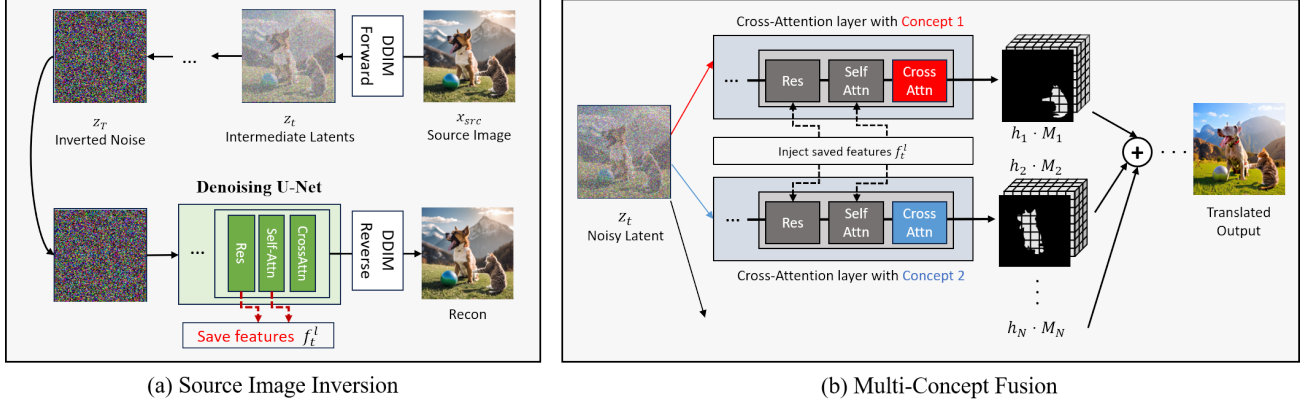


Figure 3. **Image Inversion and Multi-Concept Fusion.** (a) To extract and save the structural information of template images, we save the intermediate latent of images during the DDIM forward process. With the fully inverted noise, we extract the feature outputs from denoising U-Net during the DDIM reverse process. (b) From the noisy inverted latent, we start the multi-concept fusion generation. We denoise the noisy image with fine-tuned personalized models. After obtaining multiple cross-attention layer features, we fuse the different features from each masked region. In this step, we inject the pre-calculated self-attention and resnet features into the networks.

is expanded from the original area. To prevent confusion between overlapping regions of concepts, we kept the original dense mask only in such overlapped regions.

Step 5 : Multi-Concept Fusion. We now can generate the images with multi-concept characters as described in Figure 3(b). Since our goal is to generate images without any joint-training stage, we propose a novel sampling process which can combine the multiple single-concept personalized models in unified sampling process. Starting from inverted noisy latent z_T , we denoise the noise component from the latent. More specifically, we assume that there is a *bank of concepts* which already contains parameter sets for fine-tuned single-concept models. In practice, we select N concepts for generation, of which the weight parameters are $\theta_1, \theta_2, \dots, \theta_N$. Also, we pick one concept for background generation, which have parameters of θ_{bg} . With the selected models, we start our multi-concept fusion sampling.

One naive approach is to mix the multiple score estimation outputs similar to compositional diffusion [14]. At each time step t , the single score estimation is represented as:

$$\epsilon_{fuse} = \sum_i^N \epsilon_{\theta_i}(z_t, t, p_{+i})M_i + \epsilon_{\theta_{bg}}(z_t, t, p_{+bg})M_{bg},$$

where $\epsilon_{\theta_i}(z_t, t, p_{+i})$ is the model output from the i th concept, and M_i is the corresponding mask region for each concept. However, we found that naively mixing the different models in score estimation shows limited performance as the concepts of generated outputs are not smoothly mixed.

We address this problem by introducing multiple techniques for realistic concept-fusion:

First, we inject the pre-calculated features f_t^l to the U-net models. Since the concept-aware parameters are only related to cross-attention layers, they are not related to saved features f_t^l as they are extracted from residual and self-

attention layers. Therefore, we give the unified structural information to the entire sampling steps without deteriorating the representation of custom concepts.

Second, we found that using same text condition input to all networks yields severe artifacts and results in concept leakage problems, *i.e.* the appearance of concepts is mixed indiscriminately. Therefore, we propose a concept-aware text conditioning strategy, in which our text condition input p_{+i} contains a sentence which only includes one concept-indication modifier word. For example, if we combine two concepts of [c1] dog, [c2] cat and [bg] mountain background, our prompt construction scheme is as follows. We start from basic text prompt such as :

$$p_{base} = "A dog and a cat playing with a ball, mountain background"$$

Then we place the placeholder token in front of the each concepts for each text conditions such that:

$$p_{+1} = "A [c1] dog playing with a ball, mountain background"$$

$$p_{+2} = "A [c2] cat playing with a ball, mountain background"$$

$$p_{+bg} = "A dog and a cat playing with a ball, [bg] mountain background"$$

With the differently constructed text conditions, we can sample the concept-specific image in the targeted regions.

Third, we propose to mix the different concepts in the feature space of cross-attention layers as shown in Fig. 3(b). With the i th concept weight parameter θ_i and concept-aware prompt p_{+i} , we can extract output feature $h_i^{l,t}$ from the l th cross attention layers and timestep t . For brevity, we remove l, t as we use the feature in all layers and timesteps. With the extracted features for each concept, we can calculate mixed features such that:

$$h_{fuse} = \sum_i^N h_i M_i + h_{bg} M_{bg}.$$

We also propose a concept-free suppression method to remove the concept-free features during sampling process. Specifically, we calculate the cross attention features h_{base} from a concept-free (not fine-tuned) model $\epsilon_{\theta_{base}}$ with a basic text condition p_{base} , and extrapolate the concept-free features with the initial fused features such as:

$$h_{fuse} = (1 + \lambda) \left[\sum_i^N h_i M_i + h_{bg} M_{bg} \right] - \lambda h_{base}.$$

We then calculate the fused score estimation, such that:

$$\epsilon_{fuse} = \epsilon_{\theta}(z_t, t; h_{fuse}; f_t),$$

where h_{fuse} uses the fused features in cross attention layers, and f_t uses the pre-calculated features in self attention & residual layers.

In our model, the pre-calculated features f_t influence only the structural aspects of the image, while the fused features, represented as h_{fuse} , are exclusively concerned with concept-wise semantic information. This clear distinction ensures there is no conflict between these two components. As a result, our approach effectively accomplishes two distinct objectives: maintaining the overall structure of the template image and simultaneously altering the semantics of the objects to align with custom concepts. This dual functionality allows for a nuanced and precise manipulation of images according to specific requirements.

It is widely known that only using the conditional score estimation cannot produce proper generated outputs. Therefore, we leverage classifier-free guidance [8] to extrapolate the output from unconditional text condition $p_{\emptyset} = \emptyset$. In practice, we use the recent ‘negative’ prompt strategy instead of unconditional text condition, so that the output generated images will not contain the unwanted attributes described in the negative prompt p_{neg} . In our case, the negative-guidance score output is represented as:

$$\epsilon = \omega \cdot \epsilon_{fuse} + (1 - \omega) \cdot \epsilon_{\theta_{base}}(z_t, t, p_{neg}; f_t).$$

Implementation Details For the step 1 single-concept personalization, we adopted the official repository of Custom Diffusion [11]. We used the pre-trained Stable Diffusion V2.1(SD2.1) as our starting point for fine-tuning as the model showed improved quality. For a fair comparison, we adopted SD2.1 for all of the baseline methods. For each concept, we fine-tuned the models with 500 steps using learning rate of 1e-5. For step 2 template image generation part, we used images generated from Stable Diffusion XL with 50 sampling steps, higher resolution of 1024×1024 which takes 10 seconds for generating the image. The source image for this step can be a real images which contains the multiple objects. For step 4 mask generation, we

leveraged the pipelines from langSAM¹. For step 3 and 5, we followed the official source code of Plug-and-Play diffusion features [26]. In this stage, we also used SD2.1 as our generation backbone. We set the resolution size of generation process as 768×768, and used sampling step of 50. The entire process (from step 1 to 5) takes about 60 seconds with single RTX3090(VRAM 24GB) GPU. More sampling protocol details in the supplementary material.

4. Experimental Results

In this section, we evaluate our multi-concept fusion approach. First, we present qualitative and quantitative results that highlight our method’s effectiveness in generating multiple concepts in challenging scenarios. We then discuss our ablation, which examines the impact of different design choices. Finally, we show how our method can also be applied to edit and personalize real images.

Baselines. We compare our approach with several methods for concept personalization. We include early approaches such as Custom Diffusion [11] and Textual Inversion [3]. Moreover, we include recent approaches such as Perfusion [25] and Mix-of-show [4]. These approaches use a weight merging approach in which the model uses an optimization process to mix multiple single-concept weights into a unified set of weights. Since the Mix-of-show model uses a region-based sampling approach, we manually set the different regions for each concept for a fair comparison.

Datasets. We use diverse data sources for both quantitative and qualitative analyses. For quantitative evaluation, we select 15 distinct concepts from the Custom Concept dataset, arranged into five unique combinations. These concepts encompass a wide range of categories, including animals, humans, natural scenes, and objects. For qualitative analysis, we extend the bank of concepts with 3 animated characters concepts extracted from YouTube. The Custom Concept 101 dataset offers a wide variety of images, with each concept containing approximately 3 to 8 images. For the animated character concepts from the Blender Open Movie², we curated a collection of around 5 images per concept. The supplementary material showcases examples of all used concepts in our evaluations.

Evaluation metrics. Following [11], we assess our method against baseline approaches by measuring Text-alignment (*Text-sim*) and Image-alignment (*Image-sim*) using CLIP scores [19]. Text-alignment computes the cosine similarity between the CLIP embedding of the generated image and the CLIP embedding of the text prompt. To accurately reflect our model’s performance in generating multiple concepts, we have adapted the standard Image-alignment metric. This involves computing cosine similarity between vi-

¹<https://github.com/luca-medeiros/lang-segment-anything>

²<https://www.youtube.com/watch?v=WhWc3b3KhnY&t=52s>



Figure 4. **Qualitative Evaluation of Multi-Concept Generation.** We assess the quality of image generation by our method compared to baseline approaches, using prompts that incorporate every concept from a predefined concept bank (shown on the left). *First row:* our method successfully preserves the appearance of the target concepts while all baselines fail. *Second row:* here Mix-of-show is able to preserve the identity but struggles when the prompt includes a close interaction. *Third row:* all baseline approaches fail to generate the prompted action or to preserve the concept’s attributes; our model instead generates an image that follows the prompt while preserving the appearance of the concepts. Overall, our model generates concept-aware outputs without any concept mixing problems.

sual embeddings from designated concept regions and the embeddings of corresponding target concepts. We compute these metrics over 200 unique images generated by each model. We use 5 combinations of multiple concepts in which each combination includes more than 3 concepts. We use varied text prompts, from simple text such as “photo of dog and a cat standing, mountain background”, to complex interactions between the concepts like “photo of dog and a cat kissing, mountain background”. We report the average Text-alignment and Image-alignment scores computed over all the generated images.

4.1. Multi-Concept Generation Results

Qualitative Evaluation. We compare our method against the baselines in generating images from three-concept prompts. We include simple prompts such as “A photo of a [C1] cat and a [C2] woman standing with a [3] lighthouse background.”. We also study the generation quality for prompts involving concept interactions, for instance, “A photo of a [1] cat and a [C2] woman **hugging** with a [3] lighthouse background.”. We pick the images with the image with largest CLIP score for a fair comparison.

Figure 4 summarizes our qualitative evaluation. Most




Figure 5. **Towards More Complex Multi-Concept Generation.** We compare our method against Mix-of-show at generating images with prompts involving four challenging concepts. Mix-of-show exhibits severe problems of concept missing. Our method, instead, can successfully generate realistic concept-aware images when using a larger number of concepts.

Method	CLIP score	
	Text sim \uparrow	Image sim \uparrow
Textual Inversion	0.3423	0.7256
Custom Diffusion	0.3595	0.7875
Perfusion	0.3182	0.7563
Mix-of-show	0.3634	0.7984
Concept Weaver (ours)	0.3804	0.8124

Table 1. **Quantitative Evaluation of Multi-Concept Generation.** Our model outperforms the baselines in both CLIP scores, indicating that our outputs have better text and concept alignment.

baseline approaches [3, 11, 25] struggle to generate high-quality images, often failing to accurately capture the appearance of all target concepts and frequently mixing distinct features such as appearance, texture, or details between concepts. Mix-of-show [4] tends to generate realistic images for multi-concept prompts. However, we observe a common failure mode that mixes the concept’s appearance when the concept locations are close in space, e.g., when prompted to generate subjects that are “kissing”. In contrast, our method can successfully generate the custom concepts, even when prompted to generate interactions between these concepts, without mixing or missing concepts, therefore properly reflecting the given text prompts.

When composing more than 3 concepts, our method also outperforms the competing method of Mix-of-show as shown in Figure 5. Mix-of-show [4] requires weight mixing for multi-concept fusion, making its generated images severely deteriorated when including more concepts due to the complexity of weight optimization.

Quantitative Evaluation. Table 1 reports the CLIP scores for our method and the baseline approaches. The results showed that our method outperformed in both text-similarity and image-similarity scores which indicates that our generated outputs show better quality in both text semantic alignment and concept appearance preservation.

Human Preference Study. To further assess the perceptual quality of our generated images, we conducted a user study with 20 participants. We summarize the results in Table 2. The study was designed to capture detailed opinions along three different axes: 1) Alignment with the given text prompt (Text match), 2) Inclusion of all target con-

Method	User Study		
	Text match \uparrow	Concept match \uparrow	Realism \uparrow
Textual Inversion	2.28	1.89	2.55
Custom Diffusion	2.73	2.11	2.64
Perfusion	2.22	1.84	2.70
Mix-of-show	3.44	3.39	3.78
Concept Weaver (Ours)	4.70	4.64	4.43

Table 2. **Human Preference Study.** We assess three different axes. *Text match*: evaluates how closely the images follow a given text prompt. *Concept match*: measures the quality of preserving the appearance and attributes of target concepts. *Realism*: captures the overall quality of the generated images. We use a 5-point scale, where 1 represents “strongly disagree” and 5 “strongly agree”, and report the average across all responses.

Settings	CLIP score	
	Text sim \uparrow	Image sim \uparrow
(a) Only mask guidance	0.3140	0.7544
(b) w/o feature injection	0.3489	0.7739
(c) eps mix	0.3677	0.8023
(d) w/o concept-free suppression	0.3727	0.7936
Concept Weaver (Ours)	0.3804	0.8124

Table 3. **Ablation Study.** Quantitative comparison on ablating components of our method. We validate that each of our design choices make our model better at multi-concept generation.

cepts (Concept match), and 3) Overall quality and realism of the generated images (Realism). The participants were asked to score 20 images on each of these axis using a 5-point scale, where 1 represents “strongly disagree” and 5 “strongly agree”. More details about the protocol in the *supplementary material*. These results validate that our proposed method can generate perceptually better outputs when compared to the baseline methods, as consistently indicated by a broad range of human evaluators.

4.2. Ablation Study

We ablate our method and show a qualitative comparison between different settings in Figure 6. When we only use mask guidance similar to the approach of Mix-of-show (a), the output’s structures are severely deformed, and the image does not contain the proper concepts. (b) When we remove the feature injection, the output image again shows concept leakage and the quality is lowered. (c) When we use epsilon

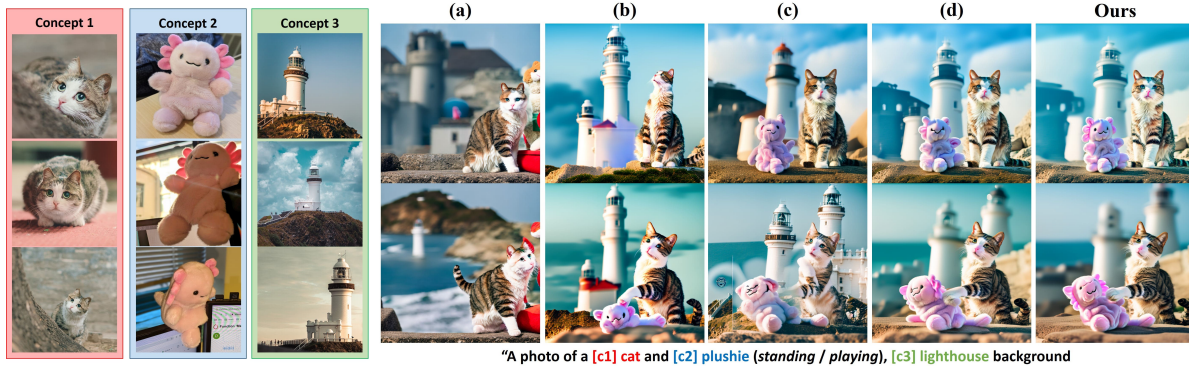


Figure 6. **Ablation Study.** (a) Results with only using mask guidance. (b) Results without using feature injection strategy. (c) Results of direct mixing on score estimation output. (d) Results without using concept-free suppression approach. (e) Ours (full).

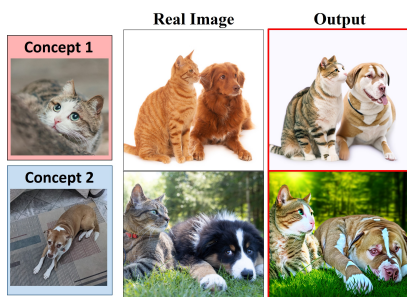


Figure 7. **Customizing Real Images.** Our method can also edit real images to inject the appearance of target concepts.

space mixing, the output image shows unwanted artifacts on the boundary area. (d) If we do not use the suppression method, the generated object does not fully reflect the concept appearance, especially for the *plushie* concept. We also show a quantitative comparison between the different settings in Table 3. We followed the same experiment protocol used in our quantitative comparison. The results validate our design choices and expose their benefits in generating images that have the highest correspondence between the text condition and the target concepts.

4.3. Applications and Potential Extensions

Customizing Real Images. Since our sampling approach starts from initial template images, we can easily extend our method into real image editing by substituting the generated template images with real ones. As shown in Figure 7, our method can edit real images with multiple custom concepts. It accurately injects the appearance and attributes of the target concepts into the existing objects in the real image.

Extension to LoRA Fine-tuning. Instead of using Custom Diffusion fine-tuning on the single-concept personalization step, we can easily adapt our approach to the more efficient scheme of Low-Rank adaptation fine-tuning. Different from the basic approach of fully fine-tuning the key and value weight W^k, W^v , we can use LoRA-based fine-tuning in which only ΔW is updated such that $W_{new} = W + \Delta W$.



Figure 8. **Extension to LoRA Fine-tuning.** Our method also supports bank of concepts trained with efficient LoRA fine-tuning.

Figure 8 illustrates that our method can easily extended to leverage the more efficient LoRA fine-tuning. We show more generated samples in Supplementary Materials.

5. Conclusion

We introduced a novel framework to generate high-fidelity images which contain multiple custom concepts. Our proposed approach fuses multiple personalized single-concept models during the sampling stage without any additional optimization process. The experimental results showed that our method outperforms state-of-the-art customization methods in multiple axes. In general, our proposed method can generate a larger number of concepts together, including complex interactions between them. We also showed that our approach can be applied to customize real images and be easily extended to efficient LoRA fine-tuning.

Acknowledgements. This research was supported by the Field oriented Technology Development Project for Customs Administration under Grant NRF2021M3I1A1097938, and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT, Ministry of Science and ICT) (No. 2022-0-00984, Development of Artificial Intelligence Technology for Personalized Plug-and-Play Explanation and Verification of Explanation, No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST))

References

- [1] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [2] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 2
- [3] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 1, 2, 5, 7
- [4] Yuchao Gu, Xintao Wang, Jay Zhangjie Wu, Yujun Shi, Yunpeng Chen, Zihan Fan, Wuyou Xiao, Rui Zhao, Shuning Chang, Weijia Wu, et al. Mix-of-show: Decentralized low-rank adaptation for multi-concept customization of diffusion models. *arXiv preprint arXiv:2305.18292*, 2023. 2, 5, 7
- [5] Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. Svdiff: Compact parameter space for diffusion fine-tuning. *arXiv preprint arXiv:2303.11305*, 2023. 2
- [6] Ligong Han, Song Wen, Qi Chen, Zhixing Zhang, Kunpeng Song, Mengwei Ren, Ruijiang Gao, Anastasis Stathopoulos, Xiaoxiao He, Yuxiao Chen, Di Liu, Qilong Zhangli, Jindong Jiang, Zhaoyang Xia, Akash Srivastava, and Dimitris Metaxas. Improving tuning-free real image editing with proximal guidance, 2023. 3
- [7] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-or. Prompt-to-prompt image editing with cross-attention control. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [8] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 5
- [9] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. 2, 3
- [10] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 3
- [11] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1931–1941, 2023. 1, 2, 3, 5, 7
- [12] Dongxu Li, Junnan Li, and Steven C. H. Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing, 2023. 2
- [13] Yuheng Li, Haotian Liu, Yangming Wen, and Yong Jae Lee. Generate anything anywhere in any scene, 2023. 2
- [14] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B. Tenenbaum. Compositional visual generation with composable diffusion models, 2023. 4
- [15] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 3
- [16] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. 2
- [17] Geon Yeong Park, Jeongsol Kim, Beomsu Kim, Sang Wan Lee, and Jong Chul Ye. Energy-based cross attention for bayesian context update in text-to-image diffusion models, 2023. 2
- [18] Can Qin, Shu Zhang, Ning Yu, Yihao Feng, Xinyi Yang, Yingbo Zhou, Huan Wang, Juan Carlos Niebles, Caiming Xiong, Silvio Savarese, et al. Unicontrol: A unified diffusion model for controllable visual generation in the wild. *arXiv preprint arXiv:2305.11147*, 2023. 2
- [19] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. 5
- [20] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [21] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2, 3
- [22] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 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*, pages 22500–22510, 2023. 1, 2
- [23] 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. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 1, 2
- [24] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [25] Yoad Tewel, Rinon Gal, Gal Chechik, and Yuval Atzmon. Key-locked rank one editing for text-to-image personalization, 2023. 1, 2, 5, 7
- [26] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. 2, 3, 5
- [27] Andrey Voynov, Qinghao Chu, Daniel Cohen-Or, and Kfir Aberman. $p+$: Extended textual conditioning in text-to-

- image generation. *arXiv preprint arXiv:2303.09522*, 2023. [2](#)
- [28] Lili Yu, Bowen Shi, Ramakanth Pasunuru, Benjamin Muller, Olga Golovneva, Tianlu Wang, Arun Babu, Binh Tang, Brian Karrer, Shelly Sheynin, et al. Scaling autoregressive multi-modal models: Pretraining and instruction tuning. *arXiv preprint arXiv:2309.02591*, 2023. [1](#), [2](#)
- [29] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks, 2017. [2](#)
- [30] Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. Inversion-based style transfer with diffusion models, 2023. [2](#)