MoMA: Multimodal LLM Adapter for Fast Personalized Image Generation

Kunpeng Song^{1,2}, Yizhe zhu¹, Bingchen Liu¹, Qing Yan¹, Ahmed Elgammal², and Xiao Yang¹

¹ ByteDance ² Rutgers University

Abstract. In this paper, we present MoMA: an open-vocabulary, trainingfree personalized image model that boasts flexible zero-shot capabilities. As foundational text-to-image models rapidly evolve, the demand for robust image-to-image translation grows. Addressing this need, MoMA specializes in subject-driven personalized image generation. Utilizing an open-source, Multimodal Large Language Model (MLLM), we train MoMA to serve a dual role as both a feature extractor and a generator. This approach effectively synergizes reference image and text prompt information to produce valuable image features, facilitating an image diffusion model. To better leverage the generated features, we further introduce a novel self-attention shortcut method that efficiently transfers image features to an image diffusion model, improving the resemblance of the target object in generated images. Remarkably, as a tuning-free plugand-play module, our model requires only a single reference image and outperforms existing methods in generating images with high detail fidelity, enhanced identity-preservation and prompt faithfulness. We commit to making our work open-source, thereby providing universal access to these advancements. Project page

Keywords: image generation · multimodal · personalization · LLM

1 Introduction

Image generation technology has seen remarkable advancements with the emergence of large-scale text-to-image diffusion models such as GLIDE [19], DALL-E 2 [25], Imagen [29], Stable Diffusion [27], eDiff-I [4]. These models enable users to generate vivid images from a diverse set of text prompts. Despite their effectiveness, textual descriptions often fall short in expressing detailed visual features, leading to the rise of image-conditioned generation works like Kandinsky [26], Stable Unclip [27, 31], which utilize images as inputs to create variations that maintain the visual components of the reference.

A natural progression in this field is subject-driven generation or image personalization. Initial efforts in this domain involve inverting input images into textual representations and employing learnable text tokens to denote target concepts. For instance, DreamBooth [28] fine-tunes the entire diffusion model to



Fig. 1: Example images generated by our open-vocabulary personalization model. With just one image of a subject (circled in blue), our model can generate textaligned, identity-preserved new images of the same subject. Our model supports both re-contextualization where the same subject is located in a new environment, as shown in green, and changing the texture of the subject itself, as shown in red.

learn a unique identifier for specific subjects. Textual-Inversion [9] inverts the input images to a unique textual embedding and learns the embedding-image mapping during finetuning. Subsequent approaches, such as Custom Diffusion [15], have optimized this process by focusing on tuning only cross-attention layers to reduce computational costs. Techniques like Low-Rank Adaptation (LoRA) [12] and SVDiff [10] further minimized trainable parameters, enhancing fine-tuning efficiency. However, these methods, regardless of their accuracy, require extensive resources for per-instance tuning and model storage, limiting their practical application.

Tuning-free approaches have gained prominence for addressing these limitations. For example, IP-Adapter [37] leverages a unique cross-attention mechanism to differentiate text and image features, facilitating the injection of reference images as visual prompts. Nevertheless, it faces notable constraints, particularly in modifying textures. Methods like InstantID [32], InstantBooth [30], and PhotoVerse [6], while maintaining identity coherence, are confined to specific domains such as human faces or cats. Recent innovations have employed transformers to integrate visual concepts with text prompts, as seen in BLIP-Diffusion [17] and KOSMOS-G [21], which extract identity information and combine it with text prompts via cross-attention. These approaches, although effective in texture modification, often result in significant detail errors in a tuning-free setting and require extra tuning for optimal outcomes with target objects.

In response to these challenges, this paper introduces a novel, open-vocabulary, and tuning-free image personalization model that excels in detail fidelity, object identity resemblance, and coherent textual prompt integration. Our model harnesses the capabilities of Multimodal Large Language Models (MLLMs) to seamlessly blend text prompts with visual features of the target object, enabling alterations in both the background context and object texture. In addition, our newly proposed self-attention shortcut significantly enhances the detail quality with minimal computational overhead.

Built upon Stable Diffusion [27] and LLaVA [?, 18], MoMA has been rigorously evaluated on various tasks with a wide array of images and dynamic prompts:

- For exact-object recontextualization tasks, it demonstrates superior detail accuracy and faithfulness to the target object across varied backgrounds.
- For texture modification tasks, our method adeptly alters the texture of target objects as dictated by text prompts while preserving unmentioned visual features.
- Our model achieves the above performance through extensive pre-training, eliminating the need for tuning at the evaluation stage.

2 Related work

2.1 Text-to-Image Diffusion Models.

Text-to-image diffusion models have become the forefront of image generation technology, garnering widespread interest for their exceptional capabilities in recent years. These models typically utilize a pre-trained language decoder to transform text prompts into latent representations that guide the diffusion process for generating or editing images. Notable models like GLIDE [19] and DisCo [33] leverage text-guided diffusion architectures and CLIP [24] guidance to enhance the fidelity and relevance of the generated images. Similarly, Stable Diffusion [27] stands out by executing the diffusion process in latent image space, thus significantly lowering computational demands. It was further advanced by Stable Diffusion XL(SDXL) [23], which introduced a larger UNet and an additional text encoder for improved textual influence on the images. Also, diffusion models have shown remarkable efficiency in capturing data distributions for image synthesis, with applications expanding by integrating transformer-based architectures [22]. The advent of text-guided image synthesis, mainly through models like Stable Diffusion, highlights significant advancements in achieving top-tier results in text-to-image synthesis tasks. Stable Diffusion, a prominent latent diffusion model, operates within a latent space defined by a pre-trained autoencoder, enabling efficient handling of semantic features and visual patterns for enhanced image synthesis.

2.2 Personalized Image Synthesis.

Recently, personalization has become an emerging factor in the vision and graphics community. Previous researchers have explored optimization-based approaches, 4 Song, K., Zhu, Y., et al.

such as Textual Inversion [9] and DreamBooth [28]. Later works found it's sufficient to just update cross-attention modules in the diffusion unet. Various parameter efficient optimization methods were then developed to further speed up tuning, such as LoRA [12] and SVDiff [10].

The drawback for these methods is they require fine-tuning for each new reference image, which is computationally expensive and time-consuming. More recent efforts attempt to get rid of per-object tuning such as [13, 30, 35], which pre-train the diffusion model on domain-specific images, such as cat and human face images. These models provide class-specific prior for generation thus require less tuning within that domain. However, they are constrained to the trained class and are not able to generalize to other subjects. Another category of works [3, 7, 17] focuses on more general open-vocabulary data. AnyDoor [7] and BreakTheScene [3] generate new images of the same object under various backgrounds but fail in changing textures. More related to our work, BLIP-Diffusion [17] uses a pre-trained transformer feature extractor and works on a wide range of subjects, however, its results contain abundant details mistakes and require few-step tuning to achieve good quality results. We posit that this is due to the lengthy information path negatively impacting the quality of image features. IP-adapter [37], as a lightweight plug-and-play model, directly injects the visual feature of the reference image into the UNet and achieves promising performance. However, the migrated image features can hardly interact with target prompts. This traps into a trade-off between prompt fidelity and image fidelity, especially when the prompt requests for drastic context changing or texture editing: higher strength results in better subject details at the cost of worse prompt faithfulness.

Nevertheless, these methods all suffer from achieving identity preservation, edibility, generalization ability and high fidelity simultaneously. Our method, however, is able to make progress in those key directions within the single input image and tuning-free domain.

3 Method

In this section, we first introduce some preliminaries about text-to-image diffusion models and multmodal LLMs. Then, we depict in detail the motivation and the design of the proposed multimodal LLM adapter.

3.1 Preliminaries.

Text-to-Image Diffusion Models Text-to-image diffusion models generate images that align with the textual description by gradually denoising a random sample drawn from a Gaussian distribution. Our work is established on text-to-image Latent Diffusion Models (LDM), which perform the diffusion process in latent space, making it more practical and computationally efficient. Given an input image x, LDM first extracts the latent feature z = E(x) with a well-trained



Fig. 2: Model structure. (1) On top-left, we adopt a generative multimodal image decoder to extract semantic features and modify them by the target prompt. These features are projected to text space and then injected into a pretrained frozen UNet with decoupled context cross-attentions as illustrated in light red. (2) On bottom-left, to further improve detail accuracy, we forward the clear reference image (t = 0) to the same UNet and extract self-attention features. These fine-grained features contain detailed information about the subject and are injected into UNet through decoupled object cross-attention layers as illustrated in orange. (3) The model is trained using a two-staged training pipeline: we first train the multimodal decoder (multimodal generative learning), then jointly optimize newly added attention modules in UNet.

encoder. During training, the noisy latent variables z_t is obtained by gradually adding noises to z for t steps, LDM optimizes the following objective:

$$\mathcal{L} = \mathbb{E}_{z_t, t, C, \varepsilon \sim \mathcal{N}(0, 1)} \left[\left\| \varepsilon - \varepsilon_{\theta}(z_t, t, C) \right\|_2^2 \right], \tag{1}$$

where C denotes the textual embedding of prompts extracted by a pre-trained CLIP text encoder.

LDM is commonly parameterized as an UNet model. We employ the pretrained Stable Diffusion as the LDM, where the UNet model has cross-attention and self-attention layers in different resolutions. Given the image features Z as query and the textual embedding C as key, the output of cross-attention Z' can be defined by the following equation:

$$Z' = Attn(Q, K, V), \tag{2}$$

where $Q = ZW_q$, $K = CW_k$, $V = CW_v$ are the query, key, and values matrices of the attention operation respectively, and W_q , W_k , W_v are the weight matrices of the trainable linear projection layers.

Multimodal LLM (MLLM) The field of natural language processing (NLP) has undergone a revolutionary shift with the emergence of large language models (LLMs). These models, distinguished by their comprehensive training across varied datasets, demonstrate extraordinary proficiency in a range of linguistic tasks.

Benefiting from robust pre-training methodologies, pioneering models, such as ChatGPT-Vision [1], Mini-GPT4 [38], CogVLM [34], and LLaVA [18], have set the stage for more sophisticated iterations in the multimodal field. These iterations notably integrate visual tasks into the LLM framework, and have proven to be exceptional in vision-related tasks, including vision-question-answering, visual-grounding, and visual-segmentation.

Among these MLLMs, LLaVA [18] stands out as an open-source Large Language and Vision Assistant, synergizing a vision encoder with an LLM for comprehensive visual and language understanding. LLaVA capitalizes on the strengths of both the pre-trained LLM and a vision transformer image encoder. This model skillfully processes images in tandem with language instructions, delivering responses in natural language, thereby bridging the gap between vision and linguistic comprehension.

3.2 Methodology

We present MoMA, a multimodal LLM adapter enhanced by fine-grained feature transfer. The overall architecture is demonstrated in Figure 2. Our method consists of three parts: (1) a generative multimodal decoder is utilized to extract image features from the reference image and edit it following the target prompt, yield the contextualized image feature; (2) in the meantime, we replace the background of the original image by white color, leaving only object pixels, leveraging the original UNet's self-attention layers to extract the object image feature; (3) finally, during the new image generation process, we injected the contextualized image features and the object image features into the UNet diffusion model with the dedicatedly trained context-cross-attention layers and object-cross-attention layers, respectively.

Multimodal Generative Image-feature Decoder We introduce a multimodal generative image-feature decoder, which actively generates target image features by combining visual information from the reference image and textual information from text prompt. Practically, we adapt a pre-trained MLLM, specifically LLaVA-7B, to serve as our generative multimodal decoder. As shown in Fig. 2 upper-left branch, given a reference image I_{ref} and its object mask M_{ref} , we get a white-background reference image by $I_{ref}^{wb} = I_{ref} * M_{ref}$.

We construct an instruction sequence as the input to MLLM: " $< f_{ref} >$ An image of < label >. Describe $< P_{tgt} >$ ", where *label* is the subject keyword (e.g. *cat*, *car*, etc.), and P_{tgt} the target prompt.

A learnable token is appended at the end of the instruction sequence. After forwarding the MLLM, the embedding corresponding to this learnable token is the output of our multimodal image-feature decoder. We call it decoder as, intuitively, it is trained to combine visual features with the target prompt in a generative manner and output an image embedding. By design, the MLLM image-feature decoder edits the background-excluded image feature of I_{ref}^{wb} following a background-included target prompt P_{tqt} that describes an entire image.



Fig. 3: Multimodal Generative Learning and iterative Self-Attention Masking

The generated image feature from the multimodal image-feature decoder is then converted into a sequence of embedding in \mathbb{R}^{768} with length N (we use N = 4 in this work) through a linear layer. Inspired by IP-Adapter [37], the embedding sequence is then integrated into the pre-trained UNet model with decoupled cross-attention as shown in Fig. 2 upper-right branch.

Self-Attention Feature Transfer To further enhance the detail faithfulness, we involve image self-attention features and apply a masking mechanism. Specifically, the same pre-trained UNet is leveraged as the self-attention feature extractor. As shown in Fig. 2 lower branch, I_{ref}^{wb} is forwarded through the diffusion UNet with t = 0 as timestep and *label* as text condition. Features at each self-attention layer are collected and transferred into the main UNet model by the adapted modules with decoupled self-attention.

The self-attention feature transfer is an effective information shortcut as the extracted feature c_i carries fine-grained details. However, directly applying it will cause interference between the backgrounds of I_{ref}^{wb} and I_{gen} . To address this issue, we present a self-attention masking procedure. Ideally, we want only the features of the foreground in I_{ref}^{wb} to be injected into the foreground of I_{gen} . The features of the background in I_{ref}^{wb} should be eliminated and the background of I_{gen} should remain unaffected by the self-attention feature transfer. We apply a masking mechanism using the reference image mask M_{ref} and the generated image mask M_{qen} . The output of our modified self-attention is:

$$Z_{new} = Attn(Q, K, V) + \lambda \cdot Attn(Q, K', V', M_{ref}) \cdot M_{gen} \cdot \beta$$
(3)

where λ is a learnable parameter. K' and V' are the key and values calculated from the extracted self-attention feature c_i by $K' = c_i W'_k$ and $V' = c_i W'_v$. Here, W'_k , W'_v are the weight matrices of the newly introduced decoupled subjectcross-attention projections. The reference image mask M_{ref} is applied inside of Attn in the form of the attention mask, and the generation mask M_{gen} is applied through an element-wise product. β is a strength scalar for additional controls.

During training, the model is optimized to reconstruct the background-included reference image I_{ref} . The white-background reference image I_{ref}^{wb} and the tar-

get image I_{ref} share the same mask, so $M_{ref} = M_{gen}$. During inference, the ground truth M_{ref} is available but the ground truth M_{gen} isn't. We use the cross-attention map corresponding to the subject *label* to approximate M_{gen} . Specifically, as shown in Fig. 3 (b), during each denoising step, the attention map of *label* from each cross-attention layer is extracted and averaged into $M_{gen(t)}$. We use it to approximate M_{gen} in the next denoising step.

Multimodal Generative Learning and Diffusion Learning. Unlike previous works that extract the image features of the subject as it is, we generate image features that are well-modified following the target text prompt. Previous works, like IP-Adapter, inject image features into the cross-attention layers of the UNet without interacting with the target prompt. This is problematic, especially when the target prompt involves texture-changing the subject. On the other hand, our multimodal image-feature decoder imagines the full image given a white-background object image and a text prompt describing the full image. which dramatically improves model performance, especially in changing subject textures. It ensures the output preserves the identity of the target object while respecting the text prompt. To achieve the best model performance, we propose a two-staged pre-training strategy.

First, we propose a **Multimodal Generative Learning Stage**, where we pre-train the multimodal image-feature decoder such that it learns to compose image features of the subject with the target prompt and output the CLIP embedding of the target image. To this end, we need to take advantage of the generative capability of the MLLM: while initially trained to generate text, we adapt it to generate image embeddings. As shown in Figure 3 (a), I_{ref}^{wb} is encoded by the MLLM vision encoder and combined with its caption P_{ref} , together with a learnable token, into a prompt instruction. This sequence is fed into the MLLM: a 15-layer transformer. The output of the learnable token is trained to match the CLIP image embedding of the original reference image I_{ref} . Once being well-trained, our MLLM will generate prompt-contexualized image embeddings. The loss function of this stage is formulated as:

$$\mathcal{L}_{\text{MLLM}} = \left\| \text{MLLM} \left(\text{CLIP} \left(I_{ref}^{wb} \right), \text{P}_{ref}, \text{Token} \right) - \text{CLIP} \left(I_{ref} \right) \right\|_{2}^{2}$$
(4)

Second, we design a **Diffusion Learning Stage** that faithfully converts the contextualized image embeddings to an image. During this stage, we freeze MLLM and pre-trained diffusion model and optimize only the decoupled subject and contextual attentions and their linear mappings. The model is trained on the OpenImage dataset, using the same training objective as shown in Eq. (1).

Classifier-free guidance (CFG) [11,37] improves diffusion generaton quality. However, we find it better to only enable it for the context-cross-attention side and not on the subject-cross-attention side. Specifically, in the second training stage, to enable CFG on the context-cross-attention side, we randomly replacing the contextualized feature with an all-zero image embedding.



Fig. 4: Zero-shot new context. Visualization of our generated samples for various images and prompts. Exact subject with a different context.

Experiments 4

Training and Implementation Detail 4.1

To train our model, we construct a dataset of 282K image/caption/image-mask triplets from the OpenImage-V7 [16] dataset. We use BLIP-2 OPT6.7B to generate captions [17] for the images, then remove human-related subjects and filter out color, shape, and texture keywords. We use the subject mask provided in OpenImage as M_{ref} . Evaluation images do not come with a mask, so we use SAM [14] to extract main objects and build masks thereafter. We use Stable Diffusion v1.5 [27] with Realistic Vision [29] checkpoint as our foundation diffusion model. We load LLaVA-7B as our MLLM decoder in stage-one training. In stage-two training, we load IP-Adapter [37] checkpoints to initialize our context cross-attention layers, and zero-initialize our object cross-attention layers. We evaluate the model using various images and prompts.

We present qualitative examples to illustrate the effectiveness of our model. In Fig. 4, the target prompts specify a novel contextual environment. Our model seamlessly generates a high-quality background while precisely situating the same object within this new setting. In Fig. 5, the prompts indicate a change in texture. Our model showcases its ability to render realistic textures in response to the textual cues, adeptly altering specified visual elements while leaving other identity aspects of the image unaffected.

Qualitative Comparison. We conducted a comparative analysis to evaluate the performance of our method against existing tuning-free open-vocabulary



Fig. 5: Zero-shot new texture. Visualization of our results with new texture, material, color and style. Our model correctly balances between prompt and image fidelity.

personalization approaches. For recontextualization, we use a variety of prompts and images to qualitatively assess our model. To ensure a fair comparison, we generated 50 random samples for each model and presented the highest-quality examples. The results, as illustrated in Fig. 6 upper panel, show our method's proficiency in detail and background quality. Particularly, the green-marble eyes and facial details in cat images is notably refined. The backgrounds generated by our model also exhibit enhanced appeal, diversity, and realism, a benefit attributed to our masked cross-attention mechanism. In texture editing, analyzed in Fig. 6 lower panel, our method consistently maintained the shape and contour of subjects while adapting them to eight different textures. This contrasts with baseline methods, which often struggled to balance the prompts and images effectively. Our results indicate a marked improvement in texture adaptation while preserving the integrity of the original subjects.

Quantitative Comparison We use the Dreambooth [28] dataset to conduct the quantitative evaluation. For a fair comparison, we generate 4 images conditioned on the image prompt for each dataset sample, resulting in 14k generated images for each method. In Tab. 1, we evaluate subject fidelity using DINO [5] and CLIP-I [24] scores, and prompt-following ability using CLIP-T [24] scores. We disaggregate the CLIP-T score into two distinct scenarios: the recontextualization task denoted by CLIP- $I_{(c)}$, and the texture editing task denoted by CLIP- $I_{(t)}$. To evaluate CLIP- $I_{(t)}$, we introduce 10 new prompts specifically tailored to texture editing. We exclude CLIP-I and DINO scores from the texture editing experiments, as they could not serve as effective metrics in this scenario:



Fig. 6: Zero-shot Qualitative Comparison. We share context editing in upper panel and texture editing in lower panel. Our results have more accurate details in context editing and better balancing between prompt and image fidelity in texture editing.

a high DINO/CLIP-I score frequently doesn't equate to high quality; rather, it indicates a failure to alter the texture as intended.

In Tab. 2, we conduct user studies on recontextualization and texture editing. Our method shows a significant performance boost across these metrics, especially in prompt-following.

Table 1: Quantitative comparisons.

 Table 2: Human Evaluations.

Methods	DINO ↑	CLIP- $T_{(c)} \uparrow$	CLIP- $I\uparrow$	$\text{CLIP-}T_{(t)}$	Metric	Ours	IP-Adapter	BLIP-Diffusion
Real Images	0.774	0.885	-	-		ouis	ii iidaptei	Din Dinabion
Textual Inversion	0.569	0.255	0.780	-	Subject Fidelity ↑	0.550	0.416	0.034
Re-Imagen	0.600	0.270	0.740	-	Prompt Fidelity \uparrow	0.633	0.351	0.016
IP-Adapter	0.612	0.330	0.793	0.319	Texture Editing ↑	0.523	0.316	0.161
BLIP-Diffusion	0.594	0.300	0.779	0.271	01			
Ours	0.618	0.348	0.803	0.335	Overall Quality \uparrow	0.596	0.368	0.0366

In Fig. 7, to visually compare model performances, we show the $\text{CLIP-}T_{(c)}$ (left) for recontextualization and $\text{CLIP-}T_{(t)}$ for texture editing experiments on 30 subjects. The left chart presents the results for the recontextualization task,



Fig. 7: Visualized CLIP-T for 30 subjects. The left chart is for the context editing task and the right is for texture editing. In each chart, one vertical axis represents one subject. Results from different models are colored differently.

where the exact subject is placed within a new context described by prompts. Our method consistently outperforms the baselines. The right chart depicts the results for the texture editing task, wherein the texture of the subject itself is altered by prompts. Both baselines experience a dramatic drop in quality. Consequently, the performance gain of our method becomes even more significant.

4.2 Ablation and Analysis

In this section, we conduct ablation studies on our proposed subject-crossattention modules and self-attention masking mechanism. The contextualized features derived from our MLLM predominantly cater to semantic understanding and the general appearance, but they inherently lack fine-grained details. To address this, we introduce a subject-cross-attention feature transfer mechanism, coupled with a masking procedure. This combination serves as a vital component for enhancing detail fidelity. The effectiveness of this approach is ablated below.

We show the effectiveness of our attention module's feature transfer ability. In Fig. 8, when $\beta = 0$, subject-cross-attention is disabled, and the model entirely relies on the multimodal decoder and its context-cross-attention injection. The generated subjects in the result images are similar to the references in shape and color, but the details are mistaken. As β increases, we observe consistent improvement in detail fidelity.

We then analyze the effectiveness of our masking mechanism. As mentioned previously, since M_{gen} is unknown in generation time, the subject-cross-attention masking itself is adapted to an iterative masking manner. As shown in Fig. 9, this design ensures the generated backgrounds are unaffected by the self-attn feature transfer and greatly improves image quality. We visualize the cross-atten map $M_{gen(t=0)}$ for the keyword *label* (in this case, dog) after the last denoising step. Notice $M_{gen(0)}$ is noisy and inaccurate when w/o masking. When it is fully activated, $M_{gen(0)}$ becomes clean and clear and the approximation to M_{gen} becomes accurate. The proposed masking technique is crucial and effective to ensure clear and diverse image backgrounds. We show quantitative results for these ablations in Tab. 3 and Tab. 4.

As further analysis, we test our model on more problem settings in addition to the previous results. As shown in Fig. 10 upper panel, our model generates



Fig. 8: Effect of self-attention feature transfer. We test various self-attn injection strengths (β) on the same fixed Z_T . Our self-attn transfer is highly effective in keeping exact subject details. As β increases from 0 to 1.2, we see results details improve consistently: the cat/dog results change from photorealism to Pixar style, and the robot/bird becomes more accurate. We show zoom-in images for the car and notice how the shape and light of the car headlight gradually become accurate as β increases.

high-quality subject-coherent images across diverse problem settings, including accessory incorporation, pose modification, and camera perspective controls. This shows it can faithfully generate subject details while allowing flexibility in prompts.

Additionally, our model is a universal adapter as we freeze the Unet in the training stage. It can generalize to the custom checkpoints fine-tuned from the same base model. In Fig. 10 lower panel, we verify this on community models from HuggingFace and Civitai [8] including Realistic Vision V4.0 [2], Anything v4 [36] and Esthetic Retro Anime [20]. These models are fine-tuned from SD v1.5. MoMA can be directly applied to these community models without modification.



Fig. 9: The effect of self-attention masking. We sample six Z_T noises and monitor their generation results with three mask strengths. As the mask strength decreases, backgrounds get more interfered with, resulting in artifacts. Iterative Masking takes care of the background and foreground separately, greatly improving background quality. The visualization of the cross-attn mask also confirms its effectiveness, as the w/ results are much less diffused than the w/o ones.

14 Song, K., Zhu, Y., et al.

Table 3: Masking Ablation.

Table 4: Subject-cross-attn Ablation.

Masking	DINO ↑	CLIP- $T_{(c)}$ 1	CLIP- $I\uparrow$	$CLIP-T_{(t)}$	β	DINO \uparrow	CLIP- $T_{(c)}$ 1	CLIP- $I\uparrow$	$CLIP-T_{(t)}$
1.0	0.618	0.348	0.803	0.335	1.0	0.618	0.348	0.803	0.335
0.5	0.603	0.322	0.797	0.231	0.5	0.591	0.348	0.788	0.338
w/o	0.575	0.298	0.788	0.322	w/o	0.583	0.349	0.782	0.319



Fig. 10: Upper panel: more results on accessory incorporation, pose and camera control. Lower panel: although trained on RealisticVision, our model can be directly applied to various community checkpoints and yields high-quality coherent results.

5 Conclusion

In conclusion, we have proposed the powerful MoMA for fast image personalization on the text-to-image diffusion model. It is tuning-free, open-vocabulary, and supports both context and texture editing. Our experiments show its superiority against existing methods. Our proposed multimodal image-feature decoder successfully harnesses the strength of MLLM for contextualized feature generation. Our masked subject-cross-attention provides a compelling feature shortcut improving detail accuracy. Additionally, as a plug-and-play module, MoMA can be integrated with community models tuned from the same base model, extending its applications to a broader area.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
- Adhik Joshi: realistic vision v40 (2024), https://huggingface.co/ stablediffusionapi/realistic-vision-v40, accessed: 2024-03-06
- Avrahami, O., Aberman, K., Fried, O., Cohen-Or, D., Lischinski, D.: Break-ascene: Extracting multiple concepts from a single image. In: SIGGRAPH Asia 2023 Conference Papers. pp. 1–12 (2023)
- Balaji, Y., Nah, S., Huang, X., Vahdat, A., Song, J., Kreis, K., Aittala, M., Aila, T., Laine, S., Catanzaro, B., et al.: ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. arXiv preprint arXiv:2211.01324 (2022)
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021)
- Chen, L., Zhao, M., Liu, Y., Ding, M., Song, Y., Wang, S., Wang, X., Yang, H., Liu, J., Du, K., et al.: Photoverse: Tuning-free image customization with text-to-image diffusion models. arXiv preprint arXiv:2309.05793 (2023)
- Chen, X., Huang, L., Liu, Y., Shen, Y., Zhao, D., Zhao, H.: Anydoor: Zero-shot object-level image customization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6593–6602 (2024)
- 8. Civitai: Civitai (2024), https://civitai.com/, accessed: 2024-03-06
- Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-Or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. arXiv preprint arXiv:2208.01618 (2022)
- Han, L., Li, Y., Zhang, H., Milanfar, P., Metaxas, D., Yang, F.: Svdiff: Compact parameter space for diffusion fine-tuning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7323–7334 (2023)
- Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
- Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021)
- Jia, X., Zhao, Y., Chan, K.C., Li, Y., Zhang, H., Gong, B., Hou, T., Wang, H., Su, Y.C.: Taming encoder for zero fine-tuning image customization with text-to-image diffusion models. arXiv preprint arXiv:2304.02642 (2023)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4015–4026 (2023)
- Kumari, N., Zhang, B., Zhang, R., Shechtman, E., Zhu, J.Y.: Multi-concept customization of text-to-image diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1931–1941 (2023)
- Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Malloci, M., Kolesnikov, A., et al.: The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. International Journal of Computer Vision **128**(7), 1956–1981 (2020)
- Li, D., Li, J., Hoi, S.: Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. Advances in Neural Information Processing Systems 36 (2024)

- 16 Song, K., Zhu, Y., et al.
- Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. Advances in neural information processing systems 36 (2024)
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
- OneRing: era-esthetic retro anime (2024), https://civitai.com/models/137781/ era-esthetic-retro-anime, accessed: 2024-03-06
- Pan, X., Dong, L., Huang, S., Peng, Z., Chen, W., Wei, F.: Kosmos-g: Generating images in context with multimodal large language models. arXiv preprint arXiv:2310.02992 (2023)
- Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4195–4205 (2023)
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical textconditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1(2), 3 (2022)
- Razzhigaev, A., Shakhmatov, A., Maltseva, A., Arkhipkin, V., Pavlov, I., Ryabov, I., Kuts, A., Panchenko, A., Kuznetsov, A., Dimitrov, D.: Kandinsky: an improved text-to-image synthesis with image prior and latent diffusion. arXiv preprint arXiv:2310.03502 (2023)
- 27. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
- Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: 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. pp. 22500–22510 (2023)
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic textto-image diffusion models with deep language understanding. Advances in neural information processing systems 35, 36479–36494 (2022)
- Shi, J., Xiong, W., Lin, Z., Jung, H.J.: Instantbooth: Personalized text-to-image generation without test-time finetuning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8543–8552 (2024)
- 31. Stable Unclip: Stable unclip (2024), https://huggingface.co/docs/diffusers/ en/api/pipelines/stable_unclip, accessed: 2024-03-06
- Wang, Q., Bai, X., Wang, H., Qin, Z., Chen, A.: Instantid: Zero-shot identitypreserving generation in seconds. arXiv preprint arXiv:2401.07519 (2024)
- Wang, T., Li, L., Lin, K., Zhai, Y., Lin, C.C., Yang, Z., Zhang, H., Liu, Z., Wang, L.: Disco: Disentangled control for realistic human dance generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9326–9336 (2024)

- 34. Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., et al.: Cogvlm: Visual expert for pretrained language models. arXiv preprint arXiv:2311.03079 (2023)
- Wei, Y., Zhang, Y., Ji, Z., Bai, J., Zhang, L., Zuo, W.: Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15943–15953 (2023)
- 36. Xyn AI: Anything v4.0 (2024), https://huggingface.co/xyn-ai/anything-v4.0, accessed: 2024-03-06
- Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. arXiv preprint arXiv:2308.06721 (2023)
- Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing visionlanguage understanding with advanced large language models. arXiv preprint arXiv:2304.10592 (2023)