SwapAnything: Enabling Arbitrary Object Swapping in Personalized Image Editing

Jing Gu^{1*} Nanxuan Zhao² Wei Xiong² Qing Liu² Zhifei Zhang² He Zhang² Jianming Zhang² HyunJoon Jung² Yilin Wang^{2**} Xin Eric Wang^{1**}

¹University of California, Santa Cruz ²Adobe https://swap-anything.github.io/



Fig. 1: *SwapAnything* results on various personalized image swapping tasks. *SwapAnything* is adept at precise, arbitrary object replacement in a source image with a personalized reference, and achieves high-fidelity swapping results without influencing any context pixels. We demonstrate its general swapping effect in single-object, multi-object, partial-object, and cross-domain swapping tasks.

Abstract. Effective editing of personal content holds a pivotal role in enabling individuals to express their creativity, weaving captivating narratives within their visual stories, and elevate the overall quality and impact of their visual content. Therefore, in this work, we introduce SwapAnything, a novel framework that can swap any objects in an image with personalized concepts given by the reference, while keeping the context unchanged. Compared with existing methods for personalized subject swapping, SwapAnything has three unique advantages: (1) precise control of arbitrary objects and parts rather than the main subject, (2) more faithful preservation of context pixels, (3) better adaptation of the personalized concept to the image. First, we propose targeted

^{*} This work was partly performed when the first author interned at Adobe.

^{**} Equal advising.

variable swapping to apply region control over latent feature maps and swap masked variables for faithful context preservation and initial semantic concept swapping. Then, we introduce appearance adaptation, to seamlessly adapt the semantic concept into the original image in terms of target location, shape, style, and content during the image generation process. Extensive results on both human and automatic evaluation demonstrate significant improvements of our approach over baseline methods on personalized swapping. Furthermore, *SwapAnything* shows its precise and faithful swapping abilities across single object, multiple objects, partial object, and cross-domain swapping tasks. *SwapAnything* also achieves great performance on text-based swapping and tasks beyond swapping such as object insertion.

1 Introduction

In today's digital age marked by the prolific creation and widespread sharing of personal content, generative models [2,7,9,30] have risen as a potent tool for self-expression, storytelling, and amplifying the impact of visual narratives. Among existing image editing methods [5, 15, 25] for content creation that empowered individuals to convey their creative instincts, weave captivating narratives into their visual stories, and enhance the quality of their visual representations, personalize content swapping [13, 14, 21], which targets at creating and compositing new images with user-specified visual concept, attracted significant interest due to its wide-ranging applications in E-commerce, entertainment, and professional editing.

Achieving arbitrary personalized content swapping necessitates a deep understanding of the visual concept inherent to both the original and replacement subjects. There are several crucial challenges. Firstly, arbitrary swapping demands significantly greater flexibility from the swapping technique compared with swapping the main subject, due to the varied nature of the content being exchanged. Secondly, an ideal swap requires flawless preservation of the surrounding context pixels, ensuring that only the designated target area undergoes modification. The third challenge lies in accurately integrating the personalized content into the target image in a harmonious manner while preserving the original poses and styles.

Existing works often fall short of addressing these challenges. Most of existing research [5,18,32,34,40] are focused on personalized image synthesis, which seeks to create new images with personalized content. Although these approaches can synthesize high-fidelity content, they cannot edit or replace the visual content in an existing image. [1,3,21,25] have tried to remove and regenerate the object via masks, they often struggle to adapt the new concept into the image. Recently studies [13,29,38] focus on object swapping and replacement by tweaking intermediate variables affecting the object's features. However, this approach lacks the precision needed for localized object swapping, resulting in limited visual qualities. Besides, those methods mainly work on main subject swapping, and they are incapable of arbitrary object swapping.

3

To address these challenges, we introduce SwapAnything, a framework that utilizes pre-trained diffusion models to streamline personalized arbitrary object swapping. Unlike previous work, our work is designed for arbitrary swapping tasks with perfect context pixel preservation and harmonious object transition. Our method begins by exploring an informative representation of the source image on a diffusion model. We found that various variables in the diffusion process especially latent features from U-net have a correspondent relation with the image. So we propose to keep the context pixels in the source image by preserving the correspondent part in those variables in the swapping process. This process is tailored to precisely swap specific areas, ensuring the preservation of other objects and the background's integrity. The object information in the source image is also selected for appearance adaptation. More specifically, location adaptation controls the location where the new concept should be swapped. Style adaptation ensures stylistic harmony between the concept object and the original image, fostering a natural and cohesive visual presentation. Additionally, scale adaptation is introduced to modulate the target object's shape, ensuring its congruence with the spatial and dimensional aspects of the source image. Last, content adaptation is crucial for smoothly generating the new concept, enabling a seamless blend that mitigates any artifacts or unnatural transitions. With these specialized adaptations, *SwapAnything* provides a heightened level of precision and refinement in the realm of object-driven image content swapping as shown in Fig. 1.

Our main contributions are: 1) We propose *SwapAnything*, a general framework for both personalized swapping and text-based swapping on single object, multiple objects, partial object, and cross-domain object. 2) We identified key variables for content preservation and proposed targeted swapping for perfect background preservation. 3) We designed a sophisticated appearance adaptation process to adapt the concept image into the source object. 4) Our approach has demonstrated exceptional performance through comprehensive qualitative assessments and quantitative analyses on swapping tasks and tasks beyond swapping such as insertion.

2 Related Work

2.1 Text-guided Image Editing

With recent progress in diffusion model, text-guided Image editing has been largely explored [1]. Image editing driven by text aims to alter an existing image following the textual guidelines provided, while ensuring certain elements or features of the original image remain unchanged. Initial efforts using GAN models [19] were restricted to specific object domains. However, diffusion-oriented techniques [11, 27, 44] have surpassed this limitation, enabling text-driven image modifications. While these techniques produce captivating outcomes, many grapple with localized edits. As a result, manual masks [25, 25, 43] are often help to delineate the editing zones. Though employing cross-attention [15] or spa-

tial attributes [38] can achieve localized edits, they face challenges with flexible changes (like altering poses) and preserving the initial image composition.

2.2 Object Driven Image Editing

In object-driven image generation, the object could be inverted into the texual space so to represented by a new token such as "*" [4, 12, 32, 34, 37]. Image editing guided by exemplars spans a vast array of uses. Much of the research in this area [16, 39, 47] falls under the umbrella of example-based image transformation tasks. These tasks leverage various sources of information, from stylized images [8, 24, 46] and layouts [17, 22, 42], to skeletons [22] and sketches or outlines [33]. Given the versatility of stylized images, the art of image style adaptation [23,45] has garnered significant interest. These methods primarily focus on establishing a detailed match between the input and reference visuals but falter when it comes to localized alterations. To facilitate localized changes, especially with flexible transformations, tools like bounding boxes and skeletons have been introduced. However, these tools often demand manual input, which can be challenging for users. A novel approach [41] frames exemplar-driven image editing as a task of filling in gaps, maintaining the context while transferring semantic elements from the reference to the original image. While DreamEdit [21] employs a step-by-step gap-filling method for object substitution, it doesn't effectively bridge the connection between the original and desired objects. In contrast, our method ensures that crucial features, like body movements and facial expressions, stay consistent. Meanwhile, Photoswap [13], CustomEdit [6] and BLIP-Diffusion [20] achieved personalized object swapping. However, it does not keep the non-object pixel background intact. While our method directly focuses on the editing area without influencing other objects and background.

3 Preliminary

Diffusion Models belong to the family of generative models that are based on stochastic processes. They generate an image by iteratively reducing noise from an initial distribution. Starting from a point of random noise denoted as \boldsymbol{z}_T , which follows a normal distribution $\boldsymbol{z}_T \sim \mathcal{N}(0, \mathbf{I})$, the diffusion model denoises each instance \boldsymbol{z}_t , thus producing \boldsymbol{z}_{t-1} . Diffusion models predict and reverse the noise at each step in the diffusion sequence to arrive at the final denoised image.

In our research, we employ the pre-trained text-to-image diffusion framework Stable Diffusion [30]. This model encodes images into a latent space and incrementally denoises the encoded latent representation. The Stable Diffusion model operates on a U-Net architecture [31], where the latent representation z_{t-1} at any given step is derived from the text prompt P and the previous latent state z_t , as indicated by the following equation:

$$\boldsymbol{z}_{t-1} = \boldsymbol{\epsilon} \boldsymbol{\theta}(\boldsymbol{z}_t, \boldsymbol{P}, t) \tag{1}$$



Fig. 2: Overview of SwapAnything on swapping a object from a source image (I_{src}) into a personalized concept $(\langle * \rangle)$ to get the target image (I_{target}) . The personalized concept is first converted into textual space to be treated as concept appearance. Meanwhile, the source image is first inverted into initial noise to obtain U-Net variables (including latent feature, attention map, and attention output). Targeted variable swapping preserves the context pixels in the source image. The appearance adaptation process then utilizes these informative variables to integrate the concept into the target image.

This U-Net includes sequence of layers that repeatedly apply self-attention and cross-attention mechanisms. In self-attention, the latent image feature z_t , is first projected into query Q_{self} , K_{self} , V_{self} , which will be used to get selfattention map A and self-attention output ϕ .

$$M = \operatorname{softmax}\left(\frac{Q_{self} \cdot K_{self}^{T}}{\sqrt{d}}\right) \qquad (2)$$
$$\phi = M \cdot V_{self}$$

Meanwhile, for cross-attention layer, the feature out of previous self-attention layer is projected into Q_{cross} , while feature embedding of textual prompt is projected into K_{cross} and V_{cross} .

$$A = \operatorname{softmax}\left(\frac{Q_{cross} \cdot K_{cross}^T}{\sqrt{d}}\right) \tag{3}$$

where A is the cross-attention map. In this work, we study the swapping of A, M, ϕ and z.



Fig. 3: Swapping process in SwapAnything. The left part shows the correspondence between latent feature z and the Generated image. The right part shows the procedure of targeted variable manipulation in the U-Net diffusion process.

4 SwapAnything

In this section, we introduce the *SwapAnything* framework that uses a diffusion model to swap on targeting area faithfully while keeping the context pixels unchanged. Fig. 2 illustrates our framework's overall structure. For source image I_{src} , we first invert it to a latent noise and then obtain the feature representations V_{src} , which will be used during the target image I_{src} generation process. In Sec. 4.1, we discuss how to preserve the non-target pixels in the source image perfectly, and how to select and transfer key information about the source image. Following this, in Sec. 4.2, we introduce the appearance adaptation pipeline that uses the key information to integrate the new concept into the source image seamlessly.

4.1 Targeted Variable Swapping

Intermediate variables in the U-Net of a diffusion model have been proven informative about the content of the generated image [3, 13, 15, 38]. They usually focus on the study of variables inside of U-net structure such as attention map, and attention output, while the output of U-Net at each diffusion step, i.e., latent image feature z is not widely explored before. We argue that the latent image feature z contains more information on image content control. The image generation process for the latent diffusion model is achieved by denoising the z to arrive at a clear representation of a high-quality image, whereas all other variables inside of U-net indirectly affect the image by impacting z. In contrast to simply swapping z like other variables, which would erase the new image's unique details and result in a mere duplication of the original image, our investigation revealed a significant correlation between the latent feature z and the produced image, including a pixel-level correspondence. In the left part of Fig. 3, we visualize the main component of the averaged latent feature z across all diffusion steps. The finding that It has a **part-to-part** correspondence with the generated images indicates the potential of localized editing by manipulating the latent feature.

7

Consequently, we consider a strategy where only the context pixels within z are altered, affecting solely the intended pixel. Although this method yields suboptimal outcomes, akin to merging two unrelated images, it prompts us to devise two remedial measures. Firstly, we suggest limiting the exchange of the latent feature to the initial stages of diffusion, allowing subsequent steps to smooth out any discordance in the latent space. Furthermore, our exploration into Unet's cross-attention map M, self-attention map A, and self-attention output ψ reveal their potential in mitigating artifacts. Swapping those facilitates the alignment of the latent features between the source and target images before the partial-swapping between them. In short, all the variables mentioned above in both the source image and target image generation process could be resized into the shape of the mask, where the mask can be utilized for the swapping process to get target variable:

$$V_{target} = g(f(V_{src}) * (1 - \mathbf{M}_{src}) + f(V_{target}) * \mathbf{M}_{src})$$
(4)

Here V includes latent feature z, and other assistant variable cross-attention map M, self-attention map A, and self-attention output ϕ . $f(\cdot)$ means the transformation process to the shape of the mask, while $g(\cdot)$ means the transformation back to the original space. For simplicity, we ignore all $f(\cdot)$ and $g(\cdot)$ in the following text. The content in the latent feature of the source image is changing as the diffusion process continues. Therefore, the location of the correspondent pixel in latent space may change over diffusion steps. A direct solution is to decode the latent feature z into an image at each step and extract the mask dynamically according to the object location in the generated image. However, we find that a changing mask usually confuses the model and leads to a less optimal performance. Therefore, we use the same high-quality mask through the diffusion process. We find that the mask could either be extracted from the source image directly using an off-the-shelf model or from the generation process as in [29,35]. Please check the Appendix for more details.

4.2 Appearance Adaptation

In this section, we introduce the appearance adaptation process that adapts the concept into the source image, which necessitates meticulous adjustments across several dimensions: location, style, scale, and content. Our framework enhances realism and coherence in image manipulation, marking a significant advancement in the field.

Location Adaptation Various intermediate variables have been proven to correlate with the final generated image. Although with great performance, it did not achieve local swap and thus the background is modified inevitably. As shown in Fig. 2, for each step, instead of directly swapping the whole variable, we conduct local swapping to only swapping the non-object position. Also, to enhance the swapping results, we further propose to conduct local swapping on the latent representation z directly. M_{src} is a 2-dimension variable containing 0 and 1. It

is the same size as of the source image and value 1 marks the swapping location. To simplify the expression, here we denote three U-Net variables attention map, attention output, and latent representation for the original image recovery process as V_{src} , denote the ones generated via target text prompt as $V_{concept}$, then we define the target variable used as V_{target}^{bg} as the background information of the target variable, which can be obtained from Eq. (5).

$$V_{target}^{bg} = V_{src} * (1 - \mathbf{M}_{src}) \tag{5}$$

The non-masked area is the swapping target area, where the variable will be generated via the target text prompt to inject the concept appearance. Location adaptation extends beyond the object swapping tasks; we also discovered its profound capability for object insertion. For further details and results, please refer to the Appendix.

Style Adaptation An ideal object swapping should keep the style unchanged. The object information in the generated variables is injected via the new concept token. Some style attributes could be already bound with the token. Therefore, solely generating the foreground information via the text prompt might lead to style inconsistency. Recently, [19, 28] found that adding such normalization layers can help improve the conditional image generation quality because such activation functions modulation. Unlike them, we employ the AdaIN (Adaptive Instance Normalization) to modulate the swapping features with spatial constraints. We follow Eq. (6) and Eq. (7) to denormalize the $V_{concept}$ with the mean and variance from V_{src} in each time step for V_{target} during the image generation process. As a result, we find that by modulating the concept feature, the generated content can adaptively follow the original style in the source image.

$$V_{concept} = MaskedAdaIN(V_{scr}, V_{concept}, M_{src})$$
(6)

$$V_{target}^{fg} = V_{concept}^{'} * M_{src} \tag{7}$$

MaskedAdaIN utilizes the mean and variance from the masked region in the AdaIN calculation. Then we have the blended feature representations for V_{target} :

$$V_{target} = V_{target}^{fg} + V_{target}^{bg} \tag{8}$$

Scale Adaptation The proportion of an object compared to its environment and other elements in the image is crucial for coherence. A swapping result with improper scaling can disturb the aesthetic balance, resulting in a disjoint appearance of the image.

Guidance from an external classifier in the inference process of Diffusion models could influence the diffusion noise to control the generated image. [10] shows that the guidance can also be used on the attention map to control the generation. Similarly, we adapt the mask guidance Eq. (9) to better align the shape between the source object and the target object.

$$\hat{\epsilon}_t = (1+s)\epsilon_{\theta}(z_t; t, y) - s\epsilon_{\theta}(z_t; t, \emptyset) + v\sigma_t \nabla_{z_t} \|M_{src} - Shape(M_{src})(k)\|_1$$
(9)

where s is the classifier-free guidance strength and v is an additional guidance weight for g. As with classifier guidance, we scale by σ_t to convert the score function to a prediction of ϵ_t . $Shape(\mathbf{M_{src}})(k)$ denotes the object shape as identified in the cross-attention layer. Here the energy function g is set as $\|\mathbf{M_{src}} - Shape(\mathbf{M_{src}})(k)\|_1$ to calculate the shape difference between the original object mask and the extracted shape of object token k in the attention layer, which indicates the deviation between the ideal shape and shape during the diffusion process.

Content Adaptation A binary mask without smoothing has a high-frequency transition at the edge — it jumps abruptly from 0 to 1. When used to merge two intermediate variables from two different diffusion processes, this can result in high-frequency artifacts at the boundary, such as jagged edges or a halo effect. Smoothing the mask transitions these high frequencies into lower frequencies, which blends the images more naturally and eliminates such artifacts. A smooth mask creates a feathering effect at the edges of the transition. This makes the merged area appear more coherent as if the two images naturally blend into each other rather than being cut off abruptly. Therefore, for the diffusion process, we specifically design two masks according to the feature of diffusion models.

Without this smoothing, the boundary between the images would be sharply defined, leading to a jarring and unnatural appearance. The Gaussian Blur softens the edges, blending the images more seamlessly. To augment this improvement, we introduce two smoothing techniques for binary masks, applied across both spatial dimensions and temporal steps. These techniques serve to refine the swapping process, mitigating artifacts and ensuring a smoother, more natural integration of the swapped regions. This results in an enriched visual output, seamlessly blending the inserted objects or object parts into the overall image composition.

Linear Boundary Interpolation: This is a process where the sharp transition between the area with 1s and the area with 0s in your binary array is made gradual. One way to achieve this is by using a convolution with a smoothing kernel (like a Gaussian kernel) that will average the values in the vicinity of each point, effectively creating a gradient at the boundary.

$$\delta(\mathbf{M}_{\mathbf{src}}) = \mathbf{M}_{\mathbf{src}} \oplus K$$

$$S = \delta(\mathbf{M}_{\mathbf{src}}) * G$$

$$S'[i,j] = \begin{cases} 1 & \text{if } M[i,j] = 1 \\ S[i,j] & \text{otherwise} \end{cases}$$
(10)

The dilation of the mask $\mathbf{M}_{\mathbf{src}}$ using the structuring element K, where denotes the dilation operation and G is the Gaussian kernel. The asterisk * denotes the convolution operation. S' is the final soft mask.

Gradual Boundary Transition: This involves generating a sequence of arrays where the value of 1 does not appear immediately but increases incrementally from 0 to 1. This can be achieved by interpolating between 0 and 1 across the sequence of arrays.

$$\mathbf{M}_{\mathbf{src}}(x,y) = \begin{cases} \mathbf{M}_{\mathbf{src}}(x,y) \cdot \frac{t}{K}, & \text{if } t \leq K \text{ and } \mathbf{M}_{\mathbf{src}}(x,y) = 1\\ \mathbf{M}_{\mathbf{src}}(x,y), & \text{otherwise} \end{cases}$$
(11)

In this equation, the value of $\mathbf{M}_{\mathbf{src}}(x, y)$ is assumed to be 1 in the center area and 0 elsewhere. For the central region, the value linearly increases from 0 to 1 over the first K steps. For the rest of the mask, the original value $\mathbf{M}_{\mathbf{src}}(x, y)$ remains unchanged.

Several prevalent backbone diffusion models, including Stable Diffusion 2.1, are restricted to processing images in a square format. Resizing images to fit a square dimension can lead to substantial content distortion, adversely affecting the editing outcomes. Nevertheless, our findings demonstrate that our method exhibits a remarkable capacity for adaptation, allowing it to **process images of any aspect ratio** without compromise. As documented in this paper, we present all images in various ratios.

5 Evaluation



Fig. 4: Qualitative comparison with different baselines. Note that those baseline methods were already equipped with some components of *SwapAnything* for precise control of the swapping region. Please check Sec. 5.1 for details.

5.1 Implementation Details

Here we introduce implementation details. In our paper, we used Stable Diffusion 2.1 as the pre-trained text-to-image diffusion model. DreamBooth [32] is used to

convert the concept into textual space. We used null-text inversion [26] based on DDIM inversion [36] to boost the accuracy. There is **no additional operation** for single-object, partial object, cross domain swapping. Multi-object swapping is achieved by conducting swapping operation on the previous swapped image. Please check Appendix for evaluation dataset details.

Baselines setting. Photoswap [13], P2P [15], PnP [38], MasaCtrl [3] are attention variable based image editing methods, which are also compatible with our proposed Masked Latent Blending in Sec. 4.1 and Location Adaptation in Sec. 4.2. Therefore we boost their performance with those additional components. Otherwise, their performance would be much worse. Please see appendix for comparisons with their original methods, and the implementation details. Also, please check appendix for its performance on object insertion.

Table 1: Human evaluation results. We show the human preference between results generated by our method and the baseline methods. SS means Object Swapping, BP means Background Preservation, and OQ means Overall Quality. SG means Object Gesture. For the baseline methods, PS means Photoswap; MC means MasaCtrl; BP means BlipDiffusion; DE means DreamEdit; CP means CopyPaste.

	Ours	$_{\rm PS}$	Tie	Ours	P2P	Tie	Ours	PnP	Tie	Ours	MC	Tie	Ours	BP	Tie	Ours	DE	Tie	Ours	CP	Tie
SS	52.3	12.5	35.2	55.3	17.9	26.8	52.3	29.9	11.5	64.3	20.0	15.7	55.3	16.5	28.2	55.1	20.1	24.8	60.1	17.3	22.6
\mathbf{SG}	44.5	34.0	21.5	49.5	32.0	18.5	55.5	33.0	11.5	65.3	18.8	15.9	41.5	39.6	18.9	44.3	18.8	36.9	54.3	20.5	25.2
BP	41.5	32.0	26.5	44.5	28.9	26.6	50.2	20.9	28.9	55.2	23.1	21.7	40.0	27.5	32.5	44.0	18.9	37.1	54.1	13.1	32.8
OQ	49.3	27.8	22.9	55.0	27.2	17.8	52.5	22.9	24.6	59.4	20.1	20.5	48.1	25.3	26.6	53.1	19.9	27.0	71.1	10.5	18.4

5.2 Single-object Swapping

We consider human evaluation to be the main quantitative performance measurement. A successful swap should keep the non-object area unchanged, change the object identity to target, and keep the gesture the same as the source object. As in Tab. 1, our model consistently outperforms baselines across all metrics. Fig. 4 shows the qualitative comparison for both human and non-human images. Thanks to the addition of our targeting variable swapping and location adaptation, all attention-manipulation based baselines also achieved perfect background preservation and some level of localized swapping result. *SwapAnything* yield a much better appearance adaptation result.

5.3 Multi-object Swapping

As is shown in Fig. 5, multi-object swapping is simply achieved via repeating single-object swapping, which further highlights its versatility and efficiency. Multi-object swapping is a natural outcome of our targeted variable swapping, whereas previous methods struggle to achieve satisfactory results. Without perfect context pixel preservation, the unwanted image modification would accumu-

late as the swapping continues. Please check the appendix for more results and comparison with baselines.



Fig. 5: Multi-object swapping results of *SwapAnything*. Our method could easily swap multiple objects via swapping one object at a time. Note that the red circle means the target object to be replaced. The same color means a pair of concept and target for object swapping.

5.4 Partial Object Swapping

As is shown in Fig. 6, *SwapAnything* achieved a great performance on swap a part of a whole object, even when the targeting area is very small. Meanwhile, all other baselines failed to achieve such results. Please check appendix for more results.



Fig. 6: Results on partial object swapping. *SwapAnything* can swap partial object that is tightly connected with other parts and adapt seamlessly to the source image. The second row is the zoom-in images of the swapping part in the first row.

5.5 Cross-domain Swapping

Fig. 7 demonstrates that *SwapAnything* can adeptly handle a range of stylized source images, successfully adapting concept objects to match the desired style within the source image while seamlessly transferring identity into the generated images. For instance, when the source image is a photo of a certain style, yet *SwapAnything* skillfully generates the same painting style featuring personalities

like "Charles Darwin" and "J. Robert Oppenheimer". Notably, all the concept images are regular, unstyled photos, underscoring the model's ability to blend different styles and identities effectively.



Fig. 7: Results on cross-domain object swapping. With a variety of source images, including graphic, black-white photos, and oil paintings, the framework seamlessly integrates concept objects taken from regular images into these diverse source images.

5.6 Text-based Swapping

As shown in Fig. 8, besides personalized swapping, our method can also do textbased swapping, swapping an object in the source image with another described in text. This can be achieved by simply replacing the personalized concept token * with a text prompt, e.g., "A photo of *new obj*".



Fig. 8: Results on text-based swapping. Besides swapping from concept image, *SwapAnything* is also capable of swapping an object described in text into the image.

5.7 Ablation Study

Fig. 9 show the effect of the components in *SwapAnything*. From the left part, we see that without latent feature swap, even with a mask and attention variable swap, the context pixel is still changed. Both latent feature and attention variable has effect of information preservation when compared with the result of no swap. The effect of adaptation and mask is presented on the right part. With style adaptation, the visual texture is closer to the source image. Without scale adaptation, the shape is not well aligned and artifacts appear in the

neck part. When without any adaptation, the generated image is much less connected the source image regarding the swapping area. When without mask, both background and targeting area are changed, which leads to a different image.



Fig. 9: Ablation study. The left part shows the effect of swapping, and the right part shows the effect of adaptation and mask.

5.8 Ethics Discussion

While SwapAnything achieves impressive performance across various visual tasks, it also raises potential ethical concerns. Like other advanced image editing technologies, SwapAnything could be misused to create deceptive content. To mitigate such risks, implementing technologies like digital watermarking can help track modifications and authenticate the sources of images, assisting users in distinguishing between genuine and altered content. Moreover, advocating for clear guidelines and regulations that govern the ethical use of image manipulation technologies is essential. These measures ensure responsible development and use, fostering a balance between innovation and ethical considerations. Here, we emphasize the importance of applications involving human face swapping. We strongly discourage using our methods on human faces without consent, and we are committed to preventing misinformation and harm to any individual or community. We enforce strict consent protocols in related tasks, requiring explicit permission from individuals before their images are used, especially in sensitive applications like face-swapping.

6 Conclusion

SwapAnything represents a notable breakthrough in the realm of object swapping. Swapping latent features and attention variables in the diffusion model ensures the retention of crucial information within the generated image. Through a targeted manipulation, we achieved a perfect background preservation. Additionally, we have introduced a sophisticated appearance adaptation process designed to seamlessly integrate the concept into the context of the source image. Consequently, SwapAnything is equipped to handle a diverse array of object swapping challenges. In the future, we plan to extend our framework to 3D/video personalized object swapping tasks.

References

- Avrahami, O., Fried, O., Lischinski, D.: Blended latent diffusion. ACM Transactions on Graphics (TOG) 42(4), 1–11 (2023)
- Blattmann, A., Rombach, R., Oktay, K., Müller, J., Ommer, B.: Retrieval-Augmented Diffusion Models. In: NIPS (2022)
- Cao, M., Wang, X., Qi, Z., Shan, Y., Qie, X., Zheng, Y.: Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In: ICCV (2023)
- 4. Chen, H., Zhang, Y., Wang, X., Duan, X., Zhou, Y., Zhu, W.: Disenbooth: Identitypreserving disentangled tuning for subject-driven text-to-image generation (2023)
- Chen, W., Hu, H., Li, Y., Rui, N., Jia, X., Chang, M.W., Cohen, W.W.: Subjectdriven text-to-image generation via apprenticeship learning. arXiv (2023)
- Choi, J., Choi, Y., Kim, Y., Kim, J., Yoon, S.: Custom-edit: Text-guided image editing with customized diffusion models. arXiv preprint arXiv:2305.15779 (2023)
- Crowson, K., Biderman, S., Kornis, D., Stander, D., Hallahan, E., Castricato, L., Raff, E.: Vqgan-clip: Open domain image generation and editing with natural language guidance. In: ECCV (2022)
- Deng, Y., Tang, F., Dong, W., Ma, C., Pan, X., Wang, L., Xu, C.: Stytr2: Image style transfer with transformers. In: CVPR (2022)
- Ding, M., Yang, Z., Hong, W., Zheng, W., Zhou, C., Yin, D., Lin, J., Zou, X., Shao, Z., Yang, H., et al.: Cogview: Mastering text-to-image generation via transformers. NIPS (2021)
- Epstein, D., Jabri, A., Poole, B., Efros, A.A., Holynski, A.: Diffusion self-guidance for controllable image generation. Advances in Neural Information Processing Systems (2023)
- Feng, W., He, X., Fu, T.J., Jampani, V., Akula, A., Narayana, P., Basu, S., Wang, X.E., Wang, W.Y.: Training-Free Structured Diffusion Guidance for Compositional Text-to-Image Synthesis. In: ICLR (2023)
- 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. In: ICLR (2023)
- Gu, J., Wang, Y., Zhao, N., Fu, T.J., Xiong, W., Liu, Q., Zhang, Z., Zhang, H., Zhang, J., Jung, H., Wang, X.E.: Photoswap: Personalized subject swapping in images (2023)
- Gu, Y., Zhou, Y., Wu, B., Yu, L., Liu, J.W., Zhao, R., Wu, J.Z., Zhang, D.J., Shou, M.Z., Tang, K.: Videoswap: Customized video subject swapping with interactive semantic point correspondence. arXiv preprint arXiv:2312.02087 (2023)
- Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-or, D.: Prompt-to-prompt image editing with cross-attention control. In: The Eleventh International Conference on Learning Representations (2022)
- Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-toimage translation. In: ECCV (2018)
- 17. Jahn, M., Rombach, R., Ommer, B.: High-resolution complex scene synthesis with transformers. arXiv (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)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019)

- 16 J. Gu, et al.
- Li, D., Li, J., Hoi, S.C.: Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. Advances in Neural Information Processing Systems (2023)
- 21. Li, T., Ku, M., Wei, C., Chen, W.: Dreamedit: Subject-driven image editing. arXiv preprint arXiv:2306.12624 (2023)
- Li, Y., Liu, H., Wu, Q., Mu, F., Yang, J., Gao, J., Li, C., Lee, Y.J.: Gligen: Open-set grounded text-to-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22511–22521 (2023)
- 23. Liao, J., Yao, Y., Yuan, L., Hua, G., Kang, S.B.: Visual attribute transfer through deep image analogy. ACM Transactions on Graphics (2017)
- 24. Liu, S., Lin, T., He, D., Li, F., Wang, M., Li, X., Sun, Z., Li, Q., Ding, E.: Adaattn: Revisit attention mechanism in arbitrary neural style transfer. In: ICCV (2021)
- Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: SDEdit: Guided image synthesis and editing with stochastic differential equations. In: International Conference on Learning Representations (2022)
- Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6038– 6047 (2023)
- Nichol, A.Q., 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. In: International Conference on Machine Learning. pp. 16784–16804. PMLR (2022)
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2337–2346 (2019)
- 29. Patashnik, O., Garibi, D., Azuri, I., Averbuch-Elor, H., Cohen-Or, D.: Localizing object-level shape variations with text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. Springer (2015)
- Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-Booth: Fine Tuning Text-to-Image Diffusion Models for Subject-Driven Generation. In: CVPR (2023)
- 33. Seo, J., Lee, G., Cho, S., Lee, J., Kim, S.: Midms: Matching interleaved diffusion models for exemplar-based image translation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 2191–2199 (2023)
- 34. Shi, J., Xiong, W., Lin, Z., Jung, H.J.: Instantbooth: Personalized text-to-image generation without test-time finetuning (2023)
- Simsar, E., Tonioni, A., Xian, Y., Hofmann, T., Tombari, F.: Lime: Localized image editing via attention regularization in diffusion models. arXiv preprint arXiv:2312.09256 (2023)
- Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2020)
- Tewel, Y., Gal, R., Chechik, G., Atzmon, Y.: Key-locked rank one editing for text-to-image personalization. In: ACM SIGGRAPH 2023 Conference Proceedings. SIGGRAPH '23 (2023)

- Tumanyan, N., Geyer, M., Bagon, S., Dekel, T.: 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. pp. 1921–1930 (2023)
- Wang, M., Yang, G.Y., Li, R., Liang, R.Z., Zhang, S.H., Hall, P.M., Hu, S.M.: Example-guided style-consistent image synthesis from semantic labeling. In: CVPR (2019)
- 40. Wang, Q., Bai, X., Wang, H., Qin, Z., Chen, A.: Instantid: Zero-shot identitypreserving generation in seconds. arXiv preprint arXiv:2401.07519 (2024)
- Yang, B., Gu, S., Zhang, B., Zhang, T., Chen, X., Sun, X., Chen, D., Wen, F.: Paint by example: Exemplar-based image editing with diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18381–18391 (2023)
- 42. Yang, Z., Wang, J., Gan, Z., Li, L., Lin, K., Wu, C., Duan, N., Liu, Z., Liu, C., Zeng, M., et al.: Reco: Region-controlled text-to-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14246–14255 (2023)
- Zeng, Y., Lin, Z., Zhang, J., Liu, Q., Collomosse, J., Kuen, J., Patel, V.M.: Scenecomposer: Any-level semantic image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22468– 22478 (2023)
- 44. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models (2023)
- Zhang, P., Zhang, B., Chen, D., Yuan, L., Wen, F.: Cross-domain correspondence learning for exemplar-based image translation. In: CVPR. pp. 5143–5153 (2020)
- 46. Zhang, Y., Huang, N., Tang, F., Huang, H., Ma, C., Dong, W., Xu, C.: Inversionbased creativity transfer with diffusion models. arXiv (2022)
- Zhou, X., Zhang, B., Zhang, T., Zhang, P., Bao, J., Chen, D., Zhang, Z., Wen, F.: Cocosnet v2: Full-resolution correspondence learning for image translation. In: CVPR (2021)