

U-VAP: User-specified Visual Appearance Personalization via Decoupled Self Augmentation

You Wu, Kean Liu, Xiaoyue Mi, Fan Tang^{1*}, Juan Cao, Jintao Li
 Institute of Computing Technology, CAS
 China

wuyou22s@ict.ac.cn, liukean22@mailsucas.ac.cn, (tangfan,mixiaoyue19s,caojuan,jtli)@ict.ac.cn




Figure 1. Given a reference image, U-VAP can personalize the user-specified visual appearances and combine them with some novel concepts. U-VAP generates images with the material (left) or color (right) extracted from the input image of a rabbit statue (middle) and achieves better accuracy and controllability in several new concepts.

Abstract

Concept personalization methods enable large text-to-image models to learn specific subjects (e.g., objects/poses/3D models) and synthesize renditions in new contexts. Given that the image references are highly biased towards visual attributes, state-of-the-art personalization models tend to overfit the whole subject and cannot disentangle visual characteristics in pixel space. In this study, we proposed a more challenging setting, namely fine-grained visual appearance personalization. Different from existing methods, we allow users to provide a sentence describing the desired attributes. A novel decoupled self-augmentation strategy is proposed to generate target-related and non-target samples to learn user-specified visual attributes. These augmented data allow for refining the model’s under-

standing of the target attribute while mitigating the impact of unrelated attributes. At the inference stage, adjustments are conducted on semantic space through the learned target and non-target embeddings to further enhance the disentanglement of target attributes. Extensive experiments on various kinds of visual attributes with SOTA personalization methods show the ability of the proposed method to mimic target visual appearance in novel contexts, thus improving the controllability and flexibility of personalization. Project page: <https://github.com/ICTMCG/U-VAP>.

1. Introduction

Recent advancements in text-to-image models [31, 34, 35, 37] allow the generation of both fantastical and realistic high-resolution visual media from mere textual prompts. To meet users’ customized generation needs, personalization

*Corresponding author: Fan Tang.

methods such as DreamBooth [36] have emerged, facilitating the learning of specific concepts from reference images and their integration with unique textual conditions. Efforts are increasingly focused on enhancing the quality and efficiency of personalized specific concept [3, 11, 12], and combination of multiple concepts [4, 19, 43, 49], significantly advancing artistic and creative applications.

Despite these strides, fine-grained visual attributes (style/layout/texture) are still challenging to extract from few-shot image inputs, which are essential for new concept generation [49]. The primary difficulty lies in the inherent entanglement of these attributes within the same pixel space and the absence of explicit external supervision. For instance, Fig. 1 demonstrates the limitations of text-conditioned guidance in DreamBooth, where the lack of explicit supervision hinders the separation of specific visual attributes. While recent works [42, 45, 46, 50] have made progress in content and style disentanglement, achieving precise attribute separation remains an elusive goal. Others rely on heuristic knowledge [44, 49] or unsupervised learning [43], lack of controllability.

In this paper, we propose a new setting for text-to-image model personalization, namely user-specified visual appearance personalization (U-VAP). To learn specific visual attributes from limited reference images in a textually controllable way, U-VAP allows users to select desired visual attributes by giving text instructions. An initial personalization model is first trained following DreamBooth [36]. However, due to the input images often sharing the same visual attributes in the given samples, the initial personalization tends to couple all the visual appearances together and ignore the users' input query.

To this end, we propose a decoupled self-augmentation strategy. Supported by the capabilities of advanced large language models [26, 28], we generate two sets of instructions according to input prompts: one set containing the target attribute with enumerating other attributes, the other vice versa. Then, the two sets of prompts are used to generate augmented samples by the initial personalization, which is used to further tune the models to couple target and non-target appearances. In this way, U-VAP promotes the personalization towards user-specified attributes and away from irrelevant attributes. Compared with SOTA personalization models, U-VAP can disentangle user-specified attributes accurately and flexibly combine them with other novel concepts. The main contributions are as follows:

- We propose U-VAP, a simple yet effective approach for achieving user-specific visual appearance personalization from limited reference images.
- We construct a decoupled self-augmentation strategy by generating target and non-target augmentations to help generate the user-specified attributes accurately.
- Experimental results demonstrate the effectiveness of U-

VAP in various attribute-aware image generation tasks, which can be combined with other customized methods in a plug-and-play manner.

2. Related Work

2.1. Personalization of diffusion models

The objective of personalization is to learn specific concepts from reference images and accurately reconstruct or freely edit them. Based on the Text-to-Image diffusion model, Textual Inversion (TI) [10] and Dreambooth [36] learn a single visual concept from several images by tuning the text encoder and entire model's parameters to bind the specific concept with a pseudo-words, respectively.

Some works [4, 8, 12, 19, 42–44, 49] have explored methods to combine multiple concepts or attributes in personalization. Custom Diffusion [19] allows for the joint training of multiple concepts or the combination of multiple fine-tuned models. With the supervision of explicit masks on a single image and cross-attention maps, Break-a-Scene [4] can learn multiple concepts separately. However, they are unable to decouple specific attributes from a concept. Different from previous concept-aware methods, Extended Textual Inversion (XTI) [44] and ProSpect [49] discovered attribute-aware properties of diffusion model from neural structure and timestep spaces, which allows for flexible learning and combination of attributes from different concepts. They cannot explicitly decouple specified attributes but rely more on empirical adjustment. Similar to attribute-aware personalization [44, 49], U-VAP aims to decouple user-specified attributes from reference concepts but enhances controllability by allowing textual guidance.

2.2. Disentanglement in Generative Models

Previous studies have demonstrated that pre-trained GANs [5, 16–18] can achieve disentanglement by traversing specific directions in their latent space [7, 13, 22, 39–41] to resulting in changes on specific attributes [1, 2, 9, 15, 30].

Recently there were some works exploring disentanglement in diffusion models [8, 20, 32, 42, 45–47]. Kwon et al. [20] and Preechakul et al. [32] explore the disentanglement with the aspect of latent representation modification for image generation and attribute transfer. Wu et al. [47] optimize the mixing weights of the source and target text embeddings for disentangled image editing. Recently, advanced image editing methods [6, 14, 24, 25] achieve more accurate attribute modification textually controllably. However, these editing methods have limitations in the diversity and flexibility of specified attribute modification. Several works [8, 42, 45, 46, 48, 50] have been proposed for extracting the specific style from reference images, which can be regarded as a coarse disentanglement only between content and style. Voynov et al. [44] and Zhang et al. [49]

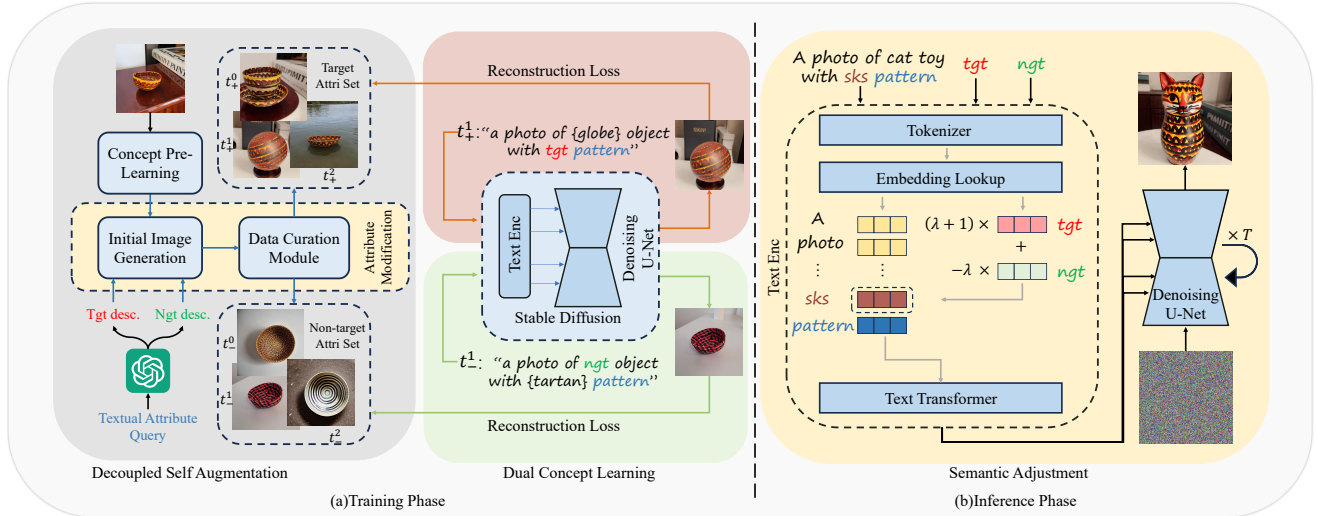


Figure 2. Pipeline of U-VAP. (a) Training: With input reference images, the initial concept-aware model is pre-trained on the entire concept. Meanwhile, given textual attribute query, U-VAP leverages GPT-3.5-turbo [26] to generate target and non-target descriptions for attribute modification. With these descriptions, U-VAP uses the initial concept-aware model to produce numerous candidate images and the data curation module filters them into target and non-target attribute sets. Subsequently, the identifiers tgt and ngt are optimized on each augmented set, which corresponds to the target and non-target attributes respectively. (b) Inference: we use semantic adjustment to correct the target embedding, further avoiding the entanglement of unwanted attributes in the generated results.

heuristically choose the embeddings obtained from neural or time spaces based on the prior knowledge of frequency characteristics of different visual attributes, which makes them difficult to distinguish attributes with similar characteristics. Vinker et al. [43] treats a single concept as a composition of multiple concepts and offers a concise but uncontrollable method for decomposition. By constructing attribute-aware samples via the self-augmentation strategy, U-VAP has the capability to controllably and flexibly personalize the specified attributes from an entire concept.

3. Method

Based on the Stable Diffusion (SD) model [35], we construct our personalization framework U-VAP on the foundation of a basic technique, DreamBooth [36]. User input a small set of reference images I_{ref} containing multiple attributes $a = \{a_{tgt}, a_{ngt}\}$, where a_{tgt} is a target attribute provided by user input prompt (e.g. “object’s structure”), and a_{ngt} is the other non-target attributes in reference images I_{ref} . Our objective is to accurately extract a particular target attribute a_{tgt} and apply it to the creation images I_{output} of new concepts using new text prompt y . It can be formalized as:

$$I_{output} = G(a_{tgt}, y, I_{ref}), \quad (1)$$

where G refers to a model used for generating images based on specific attributes. For this purpose, we extend the general personalization process and propose a framework. As shown in Fig. 2, users’ input reference images

of a concept and a textual attribute query are input data. U-VAP first learns the initial concept-aware U-Net model using DreamBooth (Sec. 3.1). Simultaneously, the decoupled self-augmentation module (Sec. 3.2) in U-VAP utilizes a large language model (LLM) to generate target and non-target attribute descriptions for attribute modification. Using these descriptions, the concept-aware model generates numerous candidate images. These are then filtered into target and non-target attribute sets. Finally, identifiers tgt and ngt are optimized for the augmented sets, corresponding to the a_{tgt} and a_{ngt} , respectively. In the inference time, semantic adjustment is employed to refine the target embedding (Sec. 3.3). This step is important in preventing the entanglement of non-target attributes in the generated images.

3.1. Preliminaries for Initial Model

In the concept pre-learning step, we leverage DreamBooth [36] to pre-learn the entire concept and obtain a concept-aware personalization model G_0 first. DreamBooth is a subject-driven personalization method, which directly fine-tunes the parameters of the diffusion model with reference images I_{ref} . It identifies the abstract categories of the subject and then constructs the text prompts using a unique identifier followed by the class name (e.g., “A sks dog”). With text prompts and reference images, DreamBooth conducts fine-tuning with the same reconstruction objective as SD. To prevent language drift [21, 23], they propose a class-specific prior preservation loss. The pa-

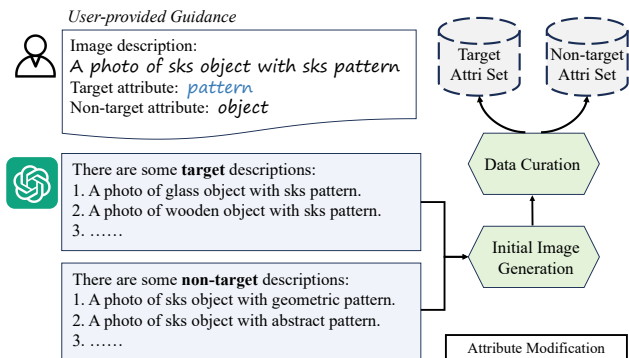


Figure 3. The workflow for decoupled self-augmentation. The “sks” represents the initial identifier learned in pre-learning step. With transformed query from guidance, an LLM [26] modifies the attribute descriptors and generates target and non-target descriptions respectively. Then the initial concept-aware model produces numerous candidate images. After data curation, U-VAP constructs attribute-aware samples for dual-concept learning.

rameters of the model are optimized simultaneously on both the reference images I_{ref} and class images generated by its original version, conditioned by text prompt and class name (e.g., A “dog”) respectively:

$$\min_{\theta} \mathbb{E}_{z,y,\varepsilon,\varepsilon',t} \left[\|\varepsilon - \varepsilon_{\theta}(z_t, t, c(y))\|_2^2 + \alpha \|\varepsilon' - \varepsilon_{\theta}(z'_t, t, c(y'))\|_2^2 \right]. \quad (2)$$

Here the latent code $z \sim \mathcal{E}(x)$ represents encoded input image I_{ref} , and \mathcal{E} is the image encoder of SD. z_t and z'_t are latent codes of user-provided image and class image by adding noise $\varepsilon \sim \mathcal{N}(0, 1)$ to z at timestep t . y is the constructed text prompt and y' is the class name for class-specific prior preservation. $c(y)$ is conditioning vector, and c is the text encoder of SD. Given noised latent z_t , the timestep t , and conditioning $c(y)$, the denoising U-Net ε_{θ} predicts and removes the added noise. α represents the weight of class-specific prior preservation loss.

3.2. Learning via Decoupled Self Augmentation

The lack of attribute-aware data is a primary obstacle in achieving precise personalization of specific attributes. Based on that, we explore a decoupled self-augmentation strategy by constructing two sets of attribute-aware samples using a large language model (LLM). Specifically, we create a target attribute text prompt set T_+ by preserving the target attributes a_{tgt} (e.g. pattern) and modifying the undesired non-target attributes a_{ngt} (e.g. object’s structure). We also construct a non-target attribute text prompt set T_- by modifying only a_{tgt} . Then we generate two sets of images I_+ and I_- by initial concept-aware model G_0 according to T_+ and T_- , respectively. Finally, we select images from I_+




Figure 4. Illustrations for semantic adjustment. By shifting semantic embedding on the adjustment direction, U-VAP can further promote the elimination of non-target attributes in the results.

and I_- that meet the criteria for constructing attribute-aware samples in the data curation module.

LLM-Powered Prompt Generation. To generate a diverse set of samples, U-VAP leverages an LLM (Specifically, GPT-3.5-turbo [26]) to automatically generate a large number of text prompts. Specifically, users first provide U-VAP with textual guidance. Then U-VAP transforms it into a query and sends it to the LLM. As shown in Fig. 3, the guidance should offer an initial description of the user’s input concepts to determine the template of text prompts. By specifying a_{tgt} and a_{ngt} , LLM can construct target descriptions T_+ and non-target descriptions T_- as text prompts respectively.

Initial Image Generation. G_0 associates the entire concept with a unique identifier *sks* after pre-learning of it. Using this identifier followed by an attribute word in the constructed text prompt (e.g., “A photo of *sks* structure backpack”), G_0 attempts to integrate this attribute onto certain new concepts (e.g., “backpack”). With these descriptions, U-VAP generates images $I_+ = G_0(T_+)$ and $I_- = G_0(T_-)$.

Data Curation. Due to the limited effectiveness of attribute modification in concept-level personalized approaches, the generated samples often fail to meet our requirements. It is necessary to filter I_+ and I_- to obtain the final attribute-aware samples $D = \{D_+, D_-\}$, represented as $D_+ = \{T_+, I_+\}$ and $D_- = \{T_-, I_-\}$. Affected by the new attribute (style or object) in the target and non-target descriptions, many images from the initial samples lose crucial information in the references. Consequently, we employ CLIP image similarity with the original image as a metric to automatically sample images with higher similarity. This criterion could at least guarantee that the filtered samples are relatively faithful to the original references, and ensures the learning process does not make the model perform worse. For introducing human intention and construct-

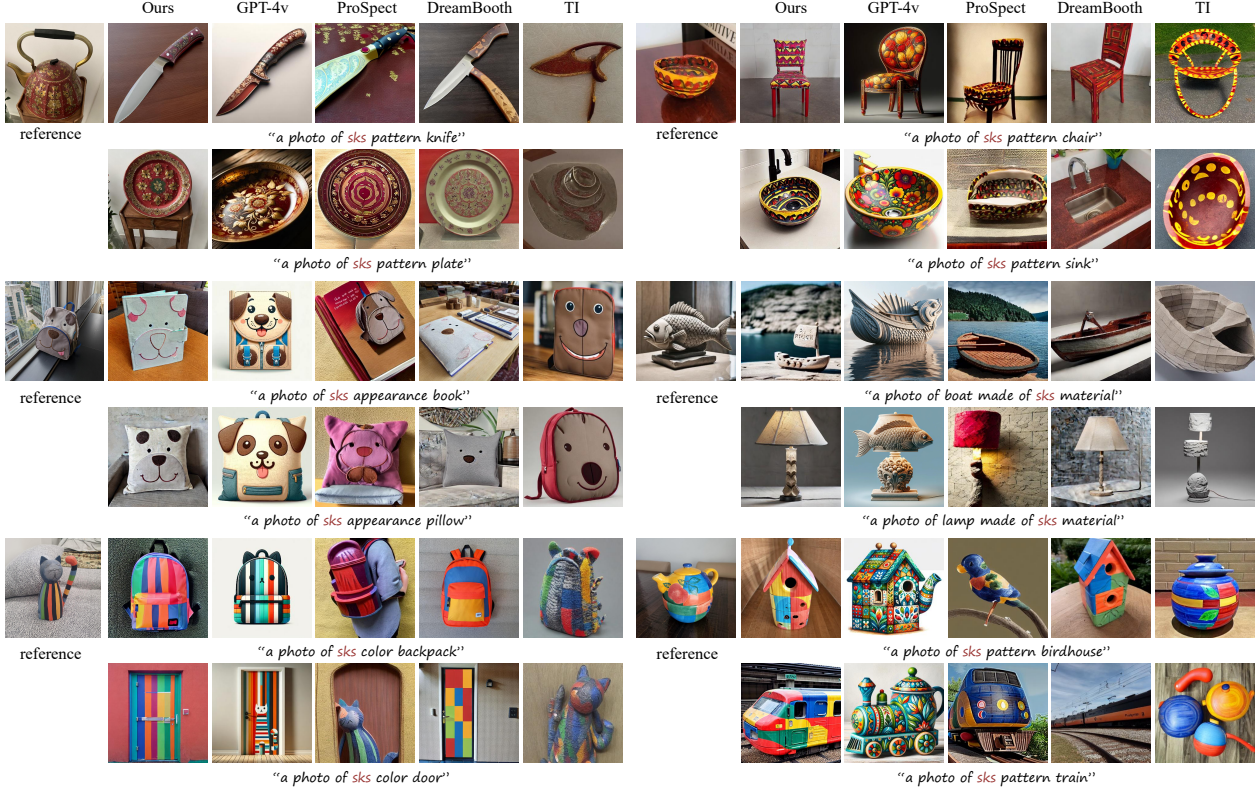


Figure 5. Qualitative comparisons. Compared with SOTA personalization methods GPT-4V, ProSpect, DreamBooth and TI, we can achieve controlled and precise generation of specific visual attributes while maintaining high visual quality.

ing better attribute-aware samples D , we strongly suggest utilizing Human Feedback for further selection after automatic filtering. Finally, we respectively sample m images for D_+ and D_- .

Dual Concept Learning. U-VAP simultaneously conducts personalization on D_+ and D_- . As illustrated in Fig. 2, the training process has two branches based on different attribute sets and learns two identifiers, t_{gt} and n_{gt} . Taking the target attribute set $D_+ = \{T_+, I_+\}$ as an example, where $T_+ = \{t_+^0, t_+^1, \dots, t_+^m\}$ and $I_+ = \{x_+^0, x_+^1, \dots, x_+^m\}$, x_+^i corresponds to a target description t_+^i and we replace the identifier sks with target identifier t_{gt} (e.g., “a photo of a globe object with a t_{gt} pattern”). When conditioned by t_+^i , the model is optimized to reconstruct x_+^i . Ultimately the identifier t_{gt} is bound to the concept described in D_+ . Due to the augmentation on the target attribute, t_{gt} is promoted to be optimized towards the semantic direction of the target attribute.

The objective of U-VAP is:

$$\min_{\theta} \mathbb{E}_{z, y, \varepsilon, \varepsilon', t} \left[\left\| \varepsilon - \varepsilon_{\theta}(z_t^+, t, c(t_+)) \right\|_2^2 + \left\| \varepsilon' - \varepsilon_{\theta}(z_t^-, t, c(t_-)) \right\|_2^2 \right], \quad (3)$$

where θ is the fine-tuning parameters of the model, z_t^+ and z_t^- are latent codes of I_+ and I_- , respectively. Different from vanilla DreamBooth, we do not consider the class-specific prior preservation loss. After dual concept learning, t_{gt} and n_{gt} independently acquired visual information related to the target attributes and non-target attributes.

3.3. Semantic Adjustment

Fig. 4 illustrates overlapping areas in semantic space for t_{gt} and n_{gt} embeddings. As a result, non-target attributes may be present in generated results. To tackle this issue, we propose a semantic adjustment during the inference phase. Starting from t_{gt} , we shift the embedding representation away from n_{gt} in the semantic space as the correction. In Fig. 2, we introduce a placeholder sks for t_{gt} to compose the text prompt p . Prior to feeding the text embedding V of p into the text transformer, the vector corresponding to sks is replaced according to:

$$v^{sks} = v^{tgt} + \lambda \overrightarrow{(v^{tgt} - v^{ngt})}, \quad (4)$$

where v^{sks} , v^{tgt} and v^{ngt} are embedding vectors of sks , t_{gt} and n_{gt} . λ represents the degree to which v^{tgt} is shifted in the direction of $v^{tgt} - v^{ngt}$. Due to the varying severity of non-target attribute entanglement at different sampling



Figure 6. Personalization of different attributes from the same reference images. Our method allows precise customization of multiple attributes (e.g., color and statue in the figure) of the reference image, as well as generalization of other new concept combinations.

noises as starting points or concepts combined, the optimal value of λ is also difficult to fix. Considering the cost of parameter adjustment, we choose to apply this operation at the inference phase.

4. Experiments

In this section, we demonstrate the effectiveness of our method in controlling personalization on various visual attributes. We conduct qualitative and quantitative experiments in Sec. 4.2, and user studies in the supplementary materials. Finally, we show sufficient ablation results in Sec. 4.3.

4.1. Experimental Setting

Training Details. Our experiments are based on Stable Diffusion v1.5 [35] with default hyperparameters. For data curation, we automatically filter 10% from approximately 200 candidate images, and 2 \sim 6 training images will be selected by human feedback (equivalent to the number of baselines). We set a learning rate of 5e-6 with 500 optimization steps at pre-learning stage. In the training phase with attribute-aware sets, we set the same learning rate with 1000 optimization steps to learn the target and non-target attributes. At the inference stage, we employed the diffusion steps in DDIM sampler of $T = 50$, and a guidance scale of 7.5. The synthesis process takes about 5 seconds, depending on the number of diffusion steps taken. All experiments are based on NVIDIA A100 with a batch size of 1.

Baselines. We choose Textual Inversion (TI) [10], DreamBooth (DB) [36], GPT-4V [29], and ProSpect [49] as our baselines to demonstrate our effectiveness in specific appearance personalization. (1) TI learns a pseudo-word for a concept within a limited number of images using an optimization-based approach. (2) DB is a widely used method that learns a unique identifier and fine-tunes the diffusion model to learn the concept from a set of im-



Figure 7. Comparison with NTI [25]+P2P [14]. The 2nd and the 5th columns are text conditions, and the 2 columns on their right are generated images by our method and NTI+P2P. “sks” is the identifier of the specific visual attribute.

ages. (3) GPT-4V is a state-of-the-art multi-modal language model, which can understand texts and images and use DALL-E 3 [27] as a tool to achieve personalization generation. (4) ProSpect learns and combines different identifiers in the expanded textual conditioning space to empirically achieve attribute personalization. Our method can combine with TI and DB in a plug-and-play manner.

Evaluation Metrics. We report quantitative results based on CLIP [33] and Inception Score (IS) [38] to evaluate both prompt fidelity, image fidelity and generation quality. Prompt fidelity is measured as the average CLIP cosine similarity between the generated image and prompt condition, which is named CLIP-T. Image fidelity is the average CLIP cosine similarity between the generated and reference images, which is named CLIP-I. Inception Score is calculated to evaluate the quality and diversity of generated images. Furthermore, we execute a study focused on user preferences to further assess these aspects in the supplementary materials.

4.2. Main Results

Quantitative Evaluation We generated 1272 images and divided target attributes into three categories: color, pattern (non-target attribute: object), and structure (like object,

shape, etc., non-target attribute: color or pattern). Then we computed prompt fidelity and Inception Score [38] to evaluate our method with TI [10], DB [36] and ProSpect [49]. As Tab. 1 shown, when learning “color” or “pattern”, our method and DB exhibit the highest CLIP-T and Inception Score (IS). DreamBooth often excessively edits towards the new concept when generating with some attributes, such as color, material, and pattern (the train in the 6th row and the 3rd column from the right in Fig. 5). Therefore, CLIP-T cannot fully illustrate DB’s editing accuracy. Excessive editing also results in increased generation quality and diversity, reflected in the highest IS score. TI [10] often results in entanglement between non-target attributes and new concepts’ attributes in images. ProSpect [49] is capable of combining attributes in some cases, but it encounters difficulties in decoupling attributes that have similar frequency characteristics. For CLIP-I, DB shows the lowest score followed by ours because it usually over-edits the non-target attributes. While Textual Inversion (TI) and ProSpect usually have higher CLIP-I, because they often face challenges in modifying objects. When learning “structure”, DB encounters the same challenges as TI and ProSpect. Our method achieves lower CLIP-I than DB, but has the highest CLIP-T because it better combines the reference structure with some new appearances.

Qualitative Evaluation As shown in Fig. 5, we compare our method with four SOTA personalization methods TI [10], DreamBooth [36], ProSpect [49], and GPT-4V [29]. To conduct a fair and unbiased evaluation, we use cases from previous papers. We augment each set of images with different texts to demonstrate the flexibility of our method. We observe that GPT-4V struggles to maintain precise target attributes. For instance, in the third-row case, GPT-4V learned a cartoon-style appearance for the dog, which was inconsistent with the realistic style of the reference image, and the details of the dog (such as the eyes and mouth) do not align with the reference. In contrast, our method successfully produces book and pillow images in a realistic style while preserving more details. ProSpect and DB show inferior performance in integrating target attributes with added semantic content. Although ProSpect maintains the target attributes, the visual quality of its generated chairs and sinks are poor. DreamBooth, on the other hand, produces a chair and sink of high quality but lacks consistency in target attributes. Our method, however, can generate these concepts with high-quality reference art patterns. TI exhibits poor performance in both consistency of target attributes and in adding semantic content (e.g., chair and sink), and sometimes overfits to the reference, leading to a lack of editability, as in the first case where it ignores textual input and generates a pot-like shape. Our approach effectively maintains target attributes while generating im-

Table 1. Quantitative evaluations. The highest results are in bold, and the second highest results are underlined.

Attributes Method	Color			Pattern			Structure		
	CLIP-I ↑	CLIP-T ↑	IS ↑	CLIP-I ↑	CLIP-T ↑	IS ↑	CLIP-I ↑	CLIP-T ↑	IS ↑
Ours	0.497	<u>0.292</u>	4.138	0.470	<u>0.270</u>	3.434	0.709	0.284	1.740
ProSpect	<u>0.581</u>	0.239	2.691	<u>0.567</u>	0.261	3.364	<u>0.773</u>	0.259	1.258
DreamBooth	0.471	0.303	<u>3.800</u>	0.467	0.299	3.793	0.837	<u>0.265</u>	1.520
TI	0.608	0.187	2.818	0.641	0.152	2.538	0.618	0.191	2.434

ages with richer semantic content.

Moreover, we demonstrate the effectiveness of our approach in personalization generation of specific appearances, by showcasing customized results for different attributes using the same reference image. As illustrated in Fig. 6, our study presents customized generation results for different attributes using the same reference image, thereby demonstrating its efficacy in controlled attribute decoupling. In Fig. 6, a colored cat statue image is used as the reference, where we have performed precise and controllable customization on both its color appearance (the first line) and statue structure (the second line). By maintaining the color appearance, our method can generate diverse objects such as multicolored backpacks, doors, coaches, and chairs. Furthermore, by preserving the “statue”, we can produce toy cats in hues of apricot, pearl, jadeite green, and gold. Those cases highlight the ability of our method to modify and control distinct attributes within an image selectively.

In Fig. 7, we compare our method with Null-text Inversion (NTI) [25] + Prompt-to-Prompt (P2P) [14], which is a diffusion-based text-guidance image editing method. NTI+P2P’s editing is considerably constrained by the basic structure and layout of the original image, resulting in edited results that do not meet the textual requirements. For example, the backpack with the color of the cat statue generated by NTI+P2P has the same structure as the reference. In contrast to NTI+P2P, our approach exhibits better editability and generates images with superior diversity.

Besides, we demonstrate that our proposed framework can also enhance the attribute personalization capabilities of Textual Inversion (TI) [10]. Unlike DreamBooth [36], TI fine-tunes the textual embedding space of SD’s text encoder to learn an identifier representing a specific concept. We can deploy TI into U-VAP by optimizing target and non-target identifiers’ embedding vectors. In Fig. 8 we compare our method with TI and Concept Decomposition [43]. As shown in Fig. 8, conditioned by “a photo of sks color chair” and “a photo of sks color dress” with the reference image, TI is almost unable to generate the color personalized images. Compared with Concept Decomposition, our method generates more accurate images of chairs and dresses, and preserves the specified color effectively. It indicates that our method can significantly improve the TI’s attribute personalization ability. Similar to Concept Decomposition [43], U-VAP is able to combine different visual attributes (as shown in the supplementary materials).



Figure 8. Comparison with TI [10] and Concept Decomposition [43]. Our method performs better in generating specified attributes.

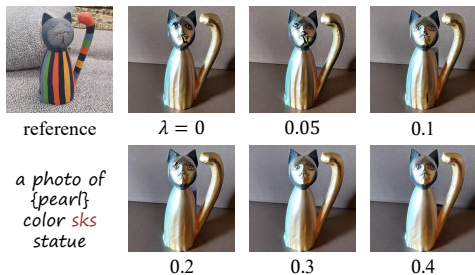


Figure 9. Ablation of λ in semantic adjustment. With λ increasing, the quality and semantic accuracy are improved.

4.3. Ablations

At the inference phase, the λ in Eq. 4 of semantic adjustment affects the accuracy and quality of generation. Fig. 9 demonstrates the influence of different values of λ when generating a new concept with the same structure as the reference. when $\lambda = 0$, semantic adjustment is not applied and leads to the appearance of undesirable non-target attributes in results, such as the stripe-like pattern. Additionally, structural details in the generated images are flawed, including distorted facial features and non-smooth edges, contributing to an overall decrease in quality. We believe that it is due to the insufficient disentanglement of target attributes learned from the target attribute set, resulting in many non-target attributes remaining. As λ increases, the embedding vector of the target identifier tgt is corrected towards a direction away from non-target attributes. When $\lambda \geq 0.3$, the visual content of unrelated attributes in the

generated images is significantly reduced. For all results without additional explanation in the paper, we set $\lambda = 0.3$.

Furthermore, the number of constructed target and non-target attribute sets also affects the personalization. More discussions are shown in the supplementary materials.

5. Conclusion

In this paper, we analyzed the reasons for the difficulty in specific visual appearance personalization. From the perspective of attribute-aware sample construction, we introduce a novel personalization framework named U-VAP. U-VAP can learn the user-specified attributes controllably and generate high-quality results with textual guidance. Extensive experiments indicate that existing works can be flexibly deployed into our framework and are improved effectively in attribute extracting. We plan to further analyze better self-augmentation methods in attribute disentanglement for making more accurate personalization in the future.

Limitations. Our method’s dependence on the capability of basic concept-aware personalization methods (such as DreamBooth [36]) results in poor performance in some attributes’ decoupling during pre-learning. For more discussions, please see supplementary materials.

Acknowledgements

This work was partly supported by the National Natural Science Foundation of China under No. 62102162, Beijing Science and Technology Plan Project under No. Z231100005923033.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4432–4441, 2019. 2
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan++: How to edit the embedded images? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8296–8305, 2020. 2
- [3] Yuval Alaluf, Elad Richardson, Gal Metzer, and Daniel Cohen-Or. A neural space-time representation for text-to-image personalization. *ACM Transactions on Graphics (TOG)*, 42(6):1–10, 2023. 2
- [4] Omri Avrahami, Kfir Aberman, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. Break-a-scene: Extracting multiple concepts from a single image. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–12, 2023. 2
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. 2
- [6] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 2
- [7] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong. Disentangled and controllable face image generation via 3d imitative-contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5154–5163, 2020. 2
- [8] Ziyi Dong, Pengxu Wei, and Liang Lin. Dreamartist: Towards controllable one-shot text-to-image generation via contrastive prompt-tuning. *arXiv preprint arXiv:2211.11337*, 2022. 2
- [9] Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 2
- [10] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 6, 7, 8
- [11] Rinon Gal, Moab Arar, Yuval Atzmon, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Encoder-based domain tuning for fast personalization of text-to-image models. *ACM Transactions on Graphics (TOG)*, 42(4):1–13, 2023. 2
- [12] Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. Svdiff: Compact parameter space for diffusion fine-tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7323–7334, 2023. 2
- [13] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. In *Advances in Neural Information Processing Systems*, pages 9841–9850. Curran Associates, Inc., 2020. 2
- [14] 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, 6, 7
- [15] Yuantian Huang, Satoshi Iizuka, Edgar Simo-Serra, and Kazuhiro Fukui. Controllable multi-domain semantic artwork synthesis. *Computational Visual Media*, 10(2):355–373, 2024. 2
- [16] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 2
- [17] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020.
- [18] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 852–863. Curran Associates, Inc., 2021. 2
- [19] 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. 2
- [20] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [21] Jason Lee, Kyunghyun Cho, and Douwe Kiela. Countering language drift via visual grounding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4385–4395, Hong Kong, China, 2019. Association for Computational Linguistics. 3
- [22] Yuheng Li, Yijun Li, Jingwan Lu, Eli Shechtman, Yong Jae Lee, and Krishna Kumar Singh. Contrastive learning for diverse disentangled foreground generation. In *European Conference on Computer Vision*, pages 334–351. Springer, 2022. 2
- [23] Yuchen Lu, Soumye Singhal, Florian Strub, Aaron Courville, and Olivier Pietquin. Countering language drift with seeded iterated learning. In *International Conference on Machine Learning*, pages 6437–6447. PMLR, 2020. 3
- [24] Daiki Miyake, Akihiro Iohara, Yu Saito, and Toshiyuki Tanaka. Negative-prompt inversion: Fast image inversion for editing with text-guided diffusion models. *arXiv preprint arXiv:2305.16807*, 2023. 2
- [25] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023. 2, 6, 7
- [26] OpenAI. GPT-3.5-turbo: Language Models for Various Applications. 2022. <https://platform.openai.com/docs/guides/gpt#gpt35-turbo>. 2, 3, 4
- [27] OpenAI. Dall-e 3 system card. Technical report, OpenAI, 2023. <https://openai.com/dall-e-3>. 6
- [28] OpenAI. Gpt-4 technical report, 2023. 2
- [29] OpenAI. GPT-4v System Card, 2023. 6, 7
- [30] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. 2
- [31] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sd-xl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 1
- [32] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10619–10629, 2022. 2
- [33] 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. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 6
- [34] 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. 1
- [35] 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, 3, 6
- [36] 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. 2, 3, 6, 7, 8
- [37] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, pages 36479–36494. Curran Associates, Inc., 2022. 1
- [38] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2016. 6, 7
- [39] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1532–1540, 2021. 2
- [40] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9243–9252, 2020.
- [41] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE transactions on pattern analysis and machine intelligence*, 44(4):2004–2018, 2020. 2
- [42] Kihyuk Sohn, Nataniel Ruiz, Kimin Lee, Daniel Castro Chin, Irina Blok, Huiwen Chang, Jarred Barber, Lu Jiang, Glenn Entis, Yuanzhen Li, et al. Styledrop: Text-to-image generation in any style. *arXiv preprint arXiv:2306.00983*, 2023. 2
- [43] Yael Vinker, Andrey Voynov, Daniel Cohen-Or, and Ariel Shamir. Concept decomposition for visual exploration and inspiration. *ACM Transactions on Graphics (TOG)*, 42(6):1–13, 2023. 2, 3, 7, 8
- [44] 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
- [45] Zhouxia Wang, Xintao Wang, Liangbin Xie, Zhongang Qi, Ying Shan, Wenping Wang, and Ping Luo. Styleadapter: A single-pass lora-free model for stylized image generation. *arXiv preprint arXiv:2309.01770*, 2023. 2
- [46] Zhizhong Wang, Lei Zhao, and Wei Xing. Stylediffusion: Controllable disentangled style transfer via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7677–7689, 2023. 2
- [47] Qiucheng Wu, Yujian Liu, Handong Zhao, Ajinkya Kale, Trung Bui, Tong Yu, Zhe Lin, Yang Zhang, and Shiyu Chang. Uncovering the disentanglement capability in text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1900–1910, 2023. 2
- [48] Yuxin Zhang, Fan Tang, Weiming Dong, Haibin Huang, Chongyang Ma, Tong-Yee Lee, and Changsheng Xu. Domain enhanced arbitrary image style transfer via contrastive learning. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–8, 2022. 2
- [49] Yuxin Zhang, Weiming Dong, Fan Tang, Nisha Huang, Haibin Huang, Chongyang Ma, Tong-Yee Lee, Oliver Deussen, and Changsheng Xu. Prospect: Prompt spectrum for attribute-aware personalization of diffusion models. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023. 2, 6, 7
- [50] Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. Inversion-based style transfer with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10146–10156, 2023. 2