

A Training-Free Style-Personalization via SVD-Based Feature Decomposition

Kyoungmin Lee*, Jihun Park*, Jongmin Gim*,
 Wonhyeok Choi, Kyumin Hwang, Jaeyeul Kim and Sunghoon Im[†]
 DGIST, Daegu, Republic of Korea

{kyoungmin, pjh2857, jongmin4422, smu06117, kyumin, jykim94, sunghoonim}@dgist.ac.kr



Figure 1. Style-personalized image generation results produced by our method. Given style reference images and text prompts, our method generates images with consistent style and diverse content.

Abstract

We present a training-free framework for style-personalized image generation that operates during inference using a scale-wise autoregressive model. Our method generates a stylized image guided by a single reference style while preserving semantic consistency and mitigating content leakage. Through a detailed step-wise analysis of the generation process, we identify a pivotal step where the dominant singular values of the internal feature encode style-

related components. Building upon this insight, we introduce two lightweight control modules: *Principal Feature Blending*, which enables precise modulation of style through SVD-based feature reconstruction, and *Structural Attention Correction*, which stabilizes structural consistency by leveraging content-guided attention correction across fine stages. Without any additional training, extensive experiments demonstrate that our method achieves competitive style fidelity and prompt fidelity compared to fine-tuned baselines, while offering faster inference and greater deployment flexibility.

*Equal contribution.

[†]Corresponding author.

1. Introduction

Text-to-Image (T2I) models [6, 16, 38, 41, 42, 45, 50] have rapidly transformed the creative landscape, enabling artists, designers, and casual users alike to generate high-quality visuals from natural language prompts. Fueled by massive image-text datasets [4, 7, 30, 46], these models now support an unprecedented range of content diversity and stylistic expression. As creative tools become increasingly democratized, users are seeking more than just visually plausible outputs—they desire personalized generation that reflects specific visual identities [27, 29, 43, 44, 57] or preferred artistic styles [18, 31, 44, 49, 64]. These emerging demands call for generation systems that are not only high-quality but also customizable, efficient, and responsive to individual preferences.

Existing solutions [1, 13, 31, 47] have made progress in this direction, often relying on fine-tuning mechanisms to encode style-specific characteristics. However, such methods typically involve training a new model instance per style, posing scalability challenges in real-world applications. Additionally, most systems are built on diffusion-based T2I architectures [26, 38, 45]. Although these models produce high-quality results, their iterative denoising process leads to slow inference, which makes them less suitable for real-time or interactive applications.

Motivated by the limitations, we propose a novel style-personalized image generation framework that combines efficiency, flexibility, and stylistic fidelity. Our method generates high-quality images guided by a single reference style image during inference without any additional training. To achieve this, we leverage a large-scale text-to-image model, specifically a scale-wise autoregressive model [16], which offers significantly faster inference compared to diffusion models while maintaining strong visual fidelity.

To better understand and maximize the capabilities of this scale-wise autoregressive model, we conduct a detailed analysis of its generation process. Our analysis reveals that a specific step in the generation process plays a crucial role in determining both content and style. At this step, the dominant singular values of the feature play a key role, as they effectively capture and separate the style-related components. Building on this insight, we develop a *Principal Feature Blending* mechanism that enables precise control over style, leveraging a specific step feature and allowing the model to faithfully reflect the stylistic traits inherent in the style features of a reference image. Additionally, we introduce a *Structural Attention Correction* strategy, which stabilizes the generation process by leveraging content-related information to preserve structural consistency. By integrating these components, our training-free framework achieves high-quality image generation, as illustrated in Fig. 1. It also achieves competitive performance in both quantitative and qualitative evaluations, while maintaining significantly

faster inference time.

In summary, our contributions include:

- We present a training-free inference framework for style-personalized image generation from a single style reference, achieving competitive results with significantly faster inference.
- We conduct a detailed step-wise analysis of the scale-wise generation process and identify a key step that governs both content and style.
- We observe that style-related components can be effectively extracted from the feature at this step through an SVD-based analysis of its dominant singular values.
- We propose two lightweight control modules—*Principal Feature Blending* for precise style modulation and *Structural Attention Correction* for stabilizing structural coherence in generation.

2. Related Work

2.1. Neural Style Transfer

Image stylization, which alters the visual style of an image, has become an active research area. A major breakthrough came with Neural Style Transfer (NST) [14], which used a pre-trained CNN (VGGNet) to separately extract content and style features. While effective, NST required costly per-image optimization. To tackle this issue, [20] proposed Adaptive Instance Normalization (AdaIN), aligning the mean and variance of the content features to those of the style features for faster style transfer. Building upon this, subsequent works [28, 33] introduced the Whitening and Coloring Transform (WCT), which aligns the full covariance structure of the features, resulting in more detailed and higher-quality stylization. With the rise of attention mechanisms in neural networks [11, 55], recent models further advanced stylization quality by utilizing attention to achieve remarkable stylization results [10, 15, 19, 32, 35, 61]. Parallel to this, vision-language models such as CLIP [40] have enabled text-driven style transfer [2, 24, 36], allowing intuitive control via natural language, without requiring explicit style images. This expands the scope of style transfer beyond traditional image-based paradigms.

2.2. Text-to-Image Generation

Recent advances in large-scale image-text datasets [4, 7, 30, 46] have greatly enhanced the ability of models to bridge the gap between natural language and visual modalities, fueling progress in conditional image synthesis. This has spurred the development of large-scale Text-to-Image (T2I) generation frameworks—including diffusion-based models [26, 38, 41, 42, 45], GAN-based approaches [22], and visual autoregressive (AR) models [6, 16, 50]—which can generate diverse, high-quality images from natural language prompts. Diffusion-based models have become the domi-

nant T2I paradigm due to their superior image quality and success in downstream tasks like style transfer and image editing [3, 9, 17, 21, 25, 39, 52, 60]. However, their high inference latency poses challenges for real-time applications. Meanwhile, visual AR models have evolved from traditional next-token prediction [12, 54] to more efficient masked token prediction [5, 6, 23]. The recent introduction of the next-scale prediction paradigm [51] has further accelerated inference without compromising output quality, establishing next-scale AR models [16, 50, 56] as promising alternatives to diffusion-based methods.

2.3. Personalized image generation

Recent advances in personalized image generation have led to methods that adapt novel visual concepts to user intent using pre-trained T2I models. These methods are generally categorized into content-oriented and style-oriented approaches. Content-oriented methods [27, 29, 43, 57] aim to capture object-specific or identity-preserving features from a small set of user-provided reference images. By fine-tuning pre-trained models or injecting learned embeddings, they generate images that maintain a high degree of fidelity to the target subject. Building on the technical foundations of these content-oriented personalization methods, recent work has extended similar principles to style-oriented personalized generation [1, 13, 18, 37, 44, 47, 49, 64]. In this paradigm, the objective shifts from preserving content identity to consistently controlling the visual style across diverse generations. Despite their effectiveness, these methods predominantly rely on diffusion-based models and often require fine-tuning, which introduces high computational costs and long inference times. In contrast, we propose a training-free, scale-wise autoregressive model that achieves fast style-personalized image generation based on a comprehensive analysis of the scale-wise autoregressive model.

3. Preliminary

Infinity Architecture. In our work, we leverage Infinity [16], a state-of-the-art T2I framework that employs the next-scale prediction paradigm introduced by [51] to generate high-fidelity, text-aligned images. During inference time, the Infinity architecture is composed of three key components: a pre-trained text encoder \mathcal{E}_T based on Flan-T5 [8], an autoregressive transformer \mathcal{M} that performs scale-wise feature prediction, and a decoder \mathcal{D} that reconstructs the final image from accumulated residual feature maps.

At each generation step $s \in \mathbf{S}$, where $\mathbf{S} = \{1, 2, \dots, S\}$ denotes the set of all generation steps, the autoregressive transformer \mathcal{M} iteratively predicts a s -th scale quantized residual feature map R_s , conditioned on the input text prompt T and the previously generated feature F_{s-1} . The process begins with initial features F_0 corresponding to the start-of-sequence (SOS) token. The prediction process is

defined as:

$$\begin{aligned} R_s &= \mathcal{M}(F_{s-1}, \mathcal{E}_T(T)) \\ &= M_{CA}(M_{SA}(Q_{s-1}, K_{s-1}, V_{s-1}), \mathcal{E}_T(T)), \forall s \in \mathbf{S}, \end{aligned} \quad (1)$$

where Q_s , K_s , and V_s are the query, key, and value at s -th generation step projected from feature F_s respectively. Here, $M_{SA}(\cdot)$ and $M_{CA}(\cdot)$ denote the self-attention and cross-attention mechanisms within the transformer.

Each predicted residual R_s is upsampled to the resolution $H \times W$ using a bilinear upsampling function $\text{up}_{H \times W}(\cdot)$, and the resulting features are accumulated across scales to form the next-step input:

$$F_s = \sum_{i=1}^s \text{up}_{H \times W}(\mathbf{R}_i), \quad R_s \in \mathbb{R}^{c \times h_s \times w_s}, \quad (2)$$

where h_s and w_s denote the spatial dimensions of the residual features at step s , c denotes the channel of the quantized feature. The final image I is produced by decoding the accumulated representation F_S at the final generation step:

$$I = \mathcal{D}(F_S). \quad (3)$$

4. Analysis of Scale-wise AR Model

(1) Step-wise analysis. We investigate the internal mechanisms of each step in the scale-wise autoregressive model’s generation process, focusing on its influence over two key visual attributes of the generated image: *content* and *style* representations. To facilitate this analysis, we construct two sets of prompt pairs $(T, \hat{T}) \in \mathbf{T}^{\text{con}} \cup \mathbf{T}^{\text{sty}}$ from 100 base prompts generated using ChatGPT, motivated by [13]:

- **Content pair set \mathbf{T}^{con} :** This set contains 100 prompt pairs $(T^{\text{con}}, \hat{T}^{\text{con}}) \in \mathbf{T}^{\text{con}}$, each formed by randomly selecting two distinct object-centric base prompts (e.g., “A photo of a donut”, “A photo of a truck”).
- **Style pair set \mathbf{T}^{sty} :** For each of the 100 base object prompts, a style pair $(T^{\text{sty}}, \hat{T}^{\text{sty}}) \in \mathbf{T}^{\text{sty}}$ was created by assigning two different colors (from a set of 10 predefined colors) to the same object, while keeping object category fixed (e.g., “A photo of a red truck” and “A photo of a green truck”).

Then, for both \mathbf{T}^{con} and \mathbf{T}^{sty} , we generate modified images \hat{I} , as shown in Fig. 2, by replacing the original text prompt T with an alternative prompt \hat{T} at a specific generation step $s \in \mathbf{S}$ in Eq. 1, while keeping all other components unchanged. A substantial change in the resulting image in response to this step-specific prompt injection indicates that the corresponding step plays a critical role in shaping certain visual attributes.

To assess the impact of each generation step, we compute the CLIP similarity [40] between alternative prompt \hat{T} and

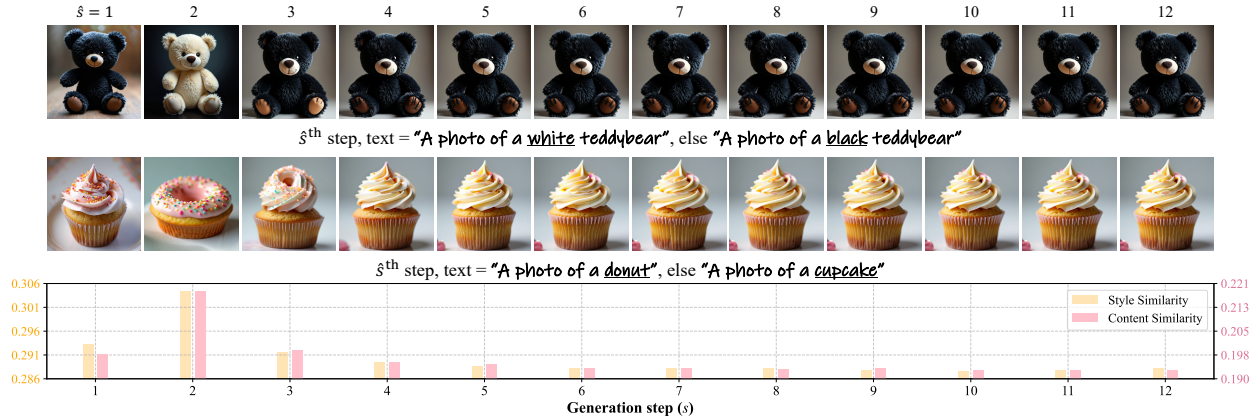


Figure 2. Step-wise prompt injection analysis. We intervene at each generation step $s \in \{1, \dots, 12\}$ by replacing the prompt only at step \hat{s} , while keeping all other steps fixed to the base prompt. **Top**: style prompt injection (“A photo of a black teddy bear” vs. “A photo of a white teddy bear”). **Middle**: content prompt injection (“A photo of a cupcake” vs. “A photo of a donut”). **Bottom**: CLIP similarity between the alternative prompt \hat{T} and the corresponding image across steps.

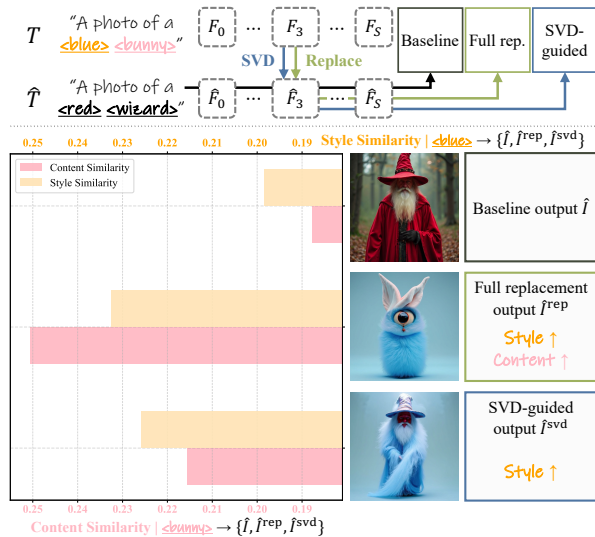


Figure 3. Key step feature analysis. Content and style similarity are measured for Baseline, Full replacement, and SVD-guided outputs using prompt pairs \mathbf{T} , with results averaged across all pairs.

its corresponding generated image as shown in the bottom row in Fig. 2. We observe that step $\hat{s} = 2$ consistently produces the highest similarity with \hat{T} across all 200 prompt pairs. This result suggests that step 2 plays a key role in shaping both content and style attributes. Consequently, the **third feature** F_3 , generated after this step, plays a crucial role in determining the final output image.

(2) Key step feature analysis. As discussed in the previous section, the third feature F_3 plays a pivotal role in shaping both content and style. Building on this observation, and supported by prior studies [34, 37] showing that early-

step features in scale-wise autoregressive models often encode strong stylistic cues, we hypothesize that the principal components of F_3 are predominantly shaped by stylistic attributes.

To analyze this, we first construct a set of 100 prompt pairs \mathbf{T} , where each pair (T, \hat{T}) differs in both object category and color (e.g., “A photo of a red truck” and “A photo of a purple cat”). For each prompt T , we apply singular value decomposition (SVD) to the third feature, yielding $F_3 = U\Sigma V^\top$. We then construct a modified diagonal matrix Σ' , obtained by zeroing out all singular values except the largest one σ_1 . Using this matrix, we reconstruct the dominant singular component as $F_3^{\text{svd}} = U\Sigma'V^\top$. The corresponding residual is defined as $F_3^{\text{res}} = F_3 - F_3^{\text{svd}}$. We perform the same decomposition for the feature \hat{F}_3 obtained from the prompt \hat{T} . For each prompt pair (T, \hat{T}) , we evaluate the effect of manipulating the dominant singular component by generating three outputs:

- **Baseline output** \hat{I} , generated using the original prompt \hat{T} without any feature manipulation.
- **Full replacement output** \hat{I}^{rep} , obtained by directly substituting the entire feature \hat{F}_3 with F_3 . ($\hat{F}_3 \leftarrow F_3$)
- **SVD-guided output** \hat{I}^{svd} , obtained by replacing only the dominant singular component of \hat{F}_3 with that of F_3 , while preserving the residual component. ($\hat{F}_3 \leftarrow F_3^{\text{svd}} + \hat{F}_3^{\text{res}}$)

As shown in Fig. 3, the full-replacement output \hat{I}^{rep} displays a substantial increase in CLIP similarity to both the object (e.g., “bunny”) and the color (e.g., “blue”) described in the substituted prompt T . In contrast, the SVD-guided output \hat{I}^{svd} shows a pronounced increase primarily in color-related CLIP similarity, while changes in object-related similarity remain much smaller. This result suggests that modifying only the dominant singular component pri-

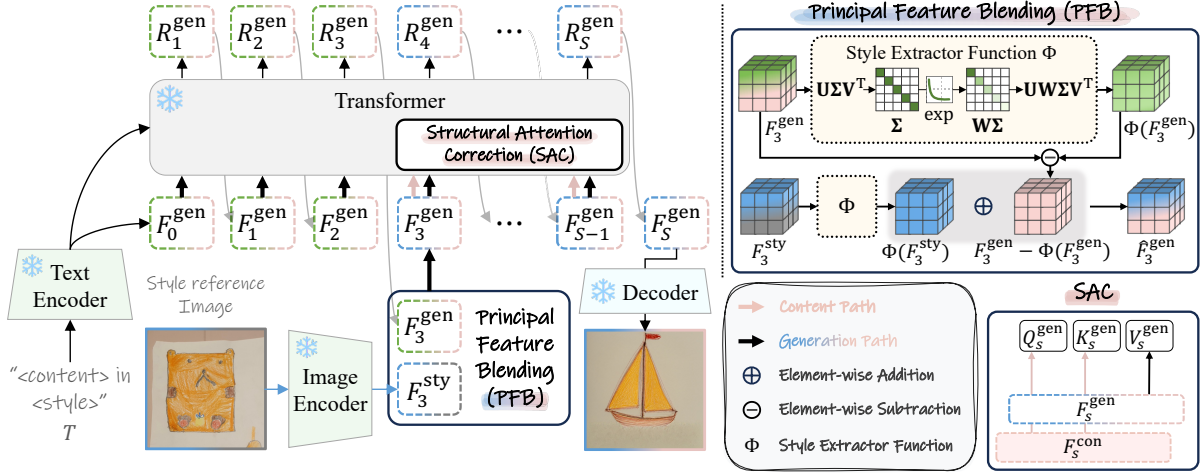


Figure 4. **Overall pipeline of our model.** The text encoder processes an identical text prompt T for both the content and generation paths, providing their embeddings to the autoregressive transformer. At stage $s = 3$, *Principal Feature Blending* is applied to extract the principal style representation from the reference style image and seamlessly integrate it into the features of the generation path. Starting from $s = 3$ (the fine stage), *Structural Attention Correction* aligns the generation path’s attention maps with those of the content path, ensuring stable and consistent structural guidance during refinement.

marily affects stylistic attributes, with minimal influence on content. Together, these observations suggest that the first principal component of the third feature F_3 predominantly captures style-related characteristics, with limited contribution from content-related information.

5. Method

5.1. Overall pipeline

In this paper, we aim to generate a style-personalized image I^{gen} by injecting a principal style feature into the final generation image, while preserving semantic consistency and suppressing content leakage.

As illustrated in Fig. 4, our method employs a dual-stream generation architecture composed of a *content path* and a *generation path*, both conditioned on the same text prompt T (“<content> in <style>”). Using an identical prompt prevents semantic mismatch between the two streams and enables consistent structural communication during inference. The *content path* operates as the standard inference branch of the pre-trained model without any modification and follows the iterative update rule in (1), producing a sequence of content features $\{F_s^{\text{con}}\}_{s=1}^S$. Its role is to provide structurally stable and semantically aligned guidance throughout the generation process. In parallel, the *generation path* follows the same update formulation but produces its own feature sequence $\{F_s^{\text{gen}}\}_{s=1}^S$, which is subsequently modulated by our proposed style-blending mechanism. This path synthesizes the final stylized output while leveraging structural cues from the content path and incorporating style information in a controlled

and targeted manner.

Building upon this dual-stream iterative process, we introduce two complementary modules—*Principal Feature Blending* (PFB) and *Structural Attention Correction* (SAC)—which operate exclusively on the generation path while leveraging cues from the other stream. PFB (Sec. 5.2) selectively injects principal style representations from the style reference image into the generation features, with a targeted intervention at the third step ($s = 3$) to prevent style-unrelated feature leakage. Following this style injection, SAC (Sec. 5.3) is applied across the subsequent steps, where it incorporates *content path* signals to stabilize structural alignment and maintain semantic consistency throughout the refinement process.

5.2. Principal Feature Blending

Our process begins by extracting the style features F_s^{sty} using a pretrained multi-scale image encoder \mathcal{E}_I from the baseline Infinity [16], as follows:

$$\{F_1^{\text{sty}}, F_2^{\text{sty}}, \dots, F_S^{\text{sty}}\} = \mathcal{E}_I(I^{\text{sty}}). \quad (4)$$

Among these multi-scale features, the step-wise analysis in Sec. 4 identifies the third feature F_3^{sty} as a pivotal representation that strongly influences both content and style during generation. Motivated by this observation, we focus on F_3^{sty} as the primary carrier of style information and use it as the basis for our style modulation mechanism.

To effectively incorporate style information while suppressing irrelevant cues from the reference image, we introduce *Principal Feature Blending*, a mechanism that selectively transfers the principal components of the style fea-

ture into the generation process. This design is grounded in two analyses from Sec. 4: (1) the **Step-wise analysis**, which pinpoints F_3 as the most influential feature for stylistic control, and (2) the **Key step feature analysis**, which reveals that the dominant singular values of F_3^{sty} encode the core stylistic characteristics. Guided by these insights, our method extracts the principal style representation from F_3^{sty} and blends it into the generation path with minimal disruption to the original content structure.

To achieve seamless blending in the feature space, we design a style extractor function Φ , which prioritizes the dominant contribution of the leading component while smoothly incorporating residual style representations. Based on the observation that the first singular value of F_3^{sty} acts as the primary carrier of style information, we enhance its influence while retaining minor contributions from the remaining components to preserve stylistic continuity. Accordingly, Φ applies exponential reweighting to the singular values based on their spectral order, gradually reducing the impact of lower components:

$$\Phi(F) \triangleq \mathbf{U}\mathbf{W}\Sigma\mathbf{V}^\top, \quad (5)$$

where $F = \mathbf{U}\Sigma\mathbf{V}^\top$,

$$\mathbf{W} = \text{diag}(\exp^{-0\cdot\alpha}, \exp^{-1\cdot\alpha}, \dots, \exp^{-(r-1)\cdot\alpha}),$$

where r denotes the rank of the feature matrix, and $\alpha > 0$ controls the exponential decay rate along the singular spectrum.

To substitute the generation path’s style representation with the one extracted from the style feature, we incorporate the refined style via Φ and update the generation feature:

$$\begin{aligned} F_3^{\text{gen}} &\leftarrow \hat{F}_3^{\text{gen}}, \\ \hat{F}_3^{\text{gen}} &= \Phi(F_3^{\text{sty}}) + (F_3^{\text{gen}} - \Phi(F_3^{\text{gen}})). \end{aligned} \quad (6)$$

This formulation preserves the original structure information of the generation path while seamlessly injecting style information derived from the reference.

5.3. Structural Attention Correction

While PFB effectively integrates style cues into the generation path, we observed that its feature-level modulation can unintentionally disturb the structural coherence of generated results, sometimes causing spatial misalignment or shape distortion. To stabilize the generation process, we leverage the attention map of the content path as a structural prior, inspired by the self-attention mechanism in diffusion-based architectures [48, 53], where the interaction between Queries and Keys preserves spatial and structural relationships. Building on this, we introduce *Structural Attention Correction (SAC)*, which aligns the attention map of the generation path with that of the content path to ensure consistent structural guidance throughout the generation process.

SAC is applied to all subsequent steps following the application of Principal Feature Blending (PFB), denoted as $\mathbf{S}_{\text{fine}} = \{3, 4, \dots, S\}$. These steps correspond to the stages where content and style representations continue to interact. Formally, SAC injects the content queries and keys at each step $s \in \mathbf{S}_{\text{fine}}$ as follows:

$$\begin{aligned} Q_s^{\text{gen}} &\leftarrow Q_s^{\text{con}}, & K_s^{\text{gen}} &\leftarrow K_s^{\text{con}}, \\ Q_s^{\text{con}} &= W_Q F_s^{\text{con}}, & K_s^{\text{con}} &= W_K F_s^{\text{con}}, \end{aligned} \quad (7)$$

Here, W_Q and W_K denote the linear projection matrices that transform input features into query and key representations in the self-attention layers. Q_s^{con} and Q_s^{gen} denote the content and generation queries at step s , and K_s^{con} and K_s^{gen} denote the corresponding keys.

6. Experiments

6.1. Implementation Details

We implement our method using a pre-trained Infinity 2B model [16] with all parameters frozen, which performs scale-wise prediction across 12 steps. The baseline employs a codebook of size 2^{32} , with quantized feature maps of resolution $64 \times 64 \times 32$. The exponential decay rate α for Principal Feature Blending is set to 1.0.

Based on our analysis, our method operates in a step-wise manner with targeted interventions across the generation process. At $s = 3$, we apply *Principal Feature Blending*, and at the fine stages ($\mathbf{S}_{\text{fine}} = \{3, 4, \dots, S\}$), we apply *Structural Attention Correction*. Generating a 1024×1024 style-personalized image takes approximately 3.58 seconds on a single NVIDIA A6000 GPU.

6.2. Evaluation Setup

Benchmark. We follow the evaluation protocol introduced in FineStyle [64], synthesizing images by combining a filtered subset of prompts from Parti [63] with 10 representative styles from the evaluation set (see Appendix for details). The Parti subset consists of 190 prompts, each describing a subject and its superclass to reduce semantic ambiguity (e.g., A cat, animals, in watercolor painting style).

Evaluation Metrics. Following FineStyle, we evaluate generated images using two CLIP-based metrics: S_{txt} (CLIP Text score) and S_{img} (CLIP Image score). S_{txt} measures the similarity between each generated image and its corresponding text prompt to assess prompt fidelity, while S_{img} measures the similarity between the generated image and a reference style image to assess style fidelity. However, a higher S_{img} does not always imply better stylization quality, as excessive similarity may result from content leakage or mode collapse, as mentioned in FineStyle. To provide a more balanced evaluation, we additionally report the harmonic mean of the two scores, denoted as S_{harmonic} (Harmonic score), which jointly reflects both prompt and style

Table 1. Quantitative comparison with state-of-the-art style-personalized image generation models. The symbols \uparrow and \downarrow indicate that higher and lower values are better, respectively. The inference time is measured as the time required to generate a single image. For tuning-based methods (*StyleDrop*, *DreamStyler*, *DB-LoRA*, and *B-LoRA*), we present the combined inference time, which accounts for both the tuning phase (given the reference style image) and the inference time required to produce a single output.

Metric	Ours	IP-Adapter	StyleAligned	StyleDrop	DreamStyler	DB-LoRA	B-LoRA	CSGO	StyleAR
Training-Free	✓	✗	✓	✗	✗	✗	✗	✗	✗
Harmonic score (S_{harmonic}) \uparrow	<u>0.437</u>	0.433	0.438	0.386	0.403	0.420	0.410	0.421	0.434
Prompt fidelity (S_{txt}) \uparrow	0.334	0.302	0.315	0.273	0.304	0.323	<u>0.324</u>	0.318	0.314
Style fidelity (S_{img}) \uparrow	0.630	0.763	<u>0.716</u>	0.657	0.599	0.602	<u>0.559</u>	0.623	0.701
Inference time (seconds) \downarrow	3.58	<u>10.13</u>	64.58	520.07	698.98	342.01	630.42	15.87	346.68




Figure 5. Qualitative comparison with state-of-the-art style-personalized image generation models.

fidelity. Formally, it is computed as:

$$S_{\text{harmonic}} = \frac{2S_{\text{txt}}S_{\text{img}}}{S_{\text{txt}} + S_{\text{img}}}. \quad (8)$$

6.3. Comparison with state-of-the-art style-personalized image generation models

To demonstrate the performance and efficiency of our model, we compare our method against eight state-of-the-art style-personalized image generation models: *StyleDrop* [49], *StyleAligned* [18], *IP-Adapter* [62], *DreamStyler* [1], *DreamBooth-LoRA* (*DB-LoRA*) [44], *B-LoRA* [13], *CSGO* [59], *StyleAR* [58].

In Tab. 1, we quantitatively compare our method with state-of-the-art baselines. While *StyleAligned* and *IP-Adapter* show relatively high style fidelity (S_{img}), they exhibit noticeably lower prompt fidelity (S_{txt}), indicating limited semantic alignment with the input text. As highlighted in [64], high S_{img} scores can be misleading due to issues like content leakage or mode collapse, where the model mimics the reference style image instead of transferring style. This effect is evident in the qualitative results shown in Fig. 5, where models with high S_{img} scores, such as *StyleAligned*

and *IP-Adapter*, frequently exhibit content leakage (first and second row). In these cases, structural details from the style reference are unintentionally transferred into the output image, leading to degraded content fidelity. These findings emphasize that high style similarity alone is insufficient to ensure faithful and semantically aligned image generation.

In contrast, *DB-LoRA* and *B-LoRA* achieve relatively high scores in prompt fidelity S_{txt} . However, they require additional fine-tuning for each new style reference, which limits scalability in practical applications. Moreover, all training-based methods suffer from long inference times, ranging from tens to hundreds of seconds per image, due to the overhead of iterative denoising or fine-tuning. Our method, by contrast, is fully training-free, up to $195\times$ faster, and achieves competitive results, making it well-suited for real-time and interactive use cases. As illustrated in Fig. 5, *DB-LoRA* and *B-LoRA* tend to better preserve the semantics of the input prompt, reinforcing that S_{txt} is a reliable indicator of semantic alignment, even when S_{img} alone may be misleading. However, despite their strong prompt adherence, both methods tend to show relatively weaker style fidelity compared to ours, suggesting that the reference style may be only partially reflected in some cases. In contrast,

our method reliably preserves both the intended content and the reference style. Despite being the fastest, it still achieves a strong balance between prompt fidelity and style fidelity, underscoring its practical advantage for high-quality, real-time style-personalized generation.

6.4. Ablation study

The quantitative results in Tab. 2 highlight how each proposed component contributes to a balance between style fidelity and prompt fidelity. To clearly demonstrate the effect of our PFB module, we compare two variants: a **direct feature replacement** strategy (REP) and our **Principal Feature Blending** (PFB). As shown in Tab. 2-(a), the baseline attains the highest prompt fidelity (S_{txt}) but exhibits limited style fidelity due to the absence of explicit style modulation. In contrast, directly replacing the style feature in (b) yields the highest style fidelity (S_{img}), but causes severe prompt degradation, indicating significant content leakage from the style reference. The SVD-based blending in (c) provides a more favorable trade-off: it mitigates the prompt-fidelity drop observed in (b) while still offering substantial improvement in style fidelity, consistent with our observation that the dominant singular component primarily captures stylistic information. Finally, the full model in (d), integrating both PFB and SAC, achieves the most balanced performance across all metrics, yielding the highest harmonic score (S_{harmonic}). This demonstrates that the proposed modules effectively complement one another, enhancing style fidelity with minimal sacrifice of prompt fidelity.

The qualitative comparison in Fig. 6 further supports these trends. In Fig. 6-(a), while the baseline produces clean and coherent images, it fails to reproduce the stylistic characteristics of the reference. Direct replacement in (b) enforces strong style transfer but also introduces unintended content elements from the reference, resulting in clear prompt mismatch. The SVD-guided variant in (c) successfully captures the intended style while retaining the target content, though its prompt adherence is still weaker than the baseline. In contrast, the full model in (d) preserves the style of the reference and simultaneously generates images that align closely with the prompt, achieving the most balanced and desirable output—consistent with the quantitative trends observed above.

6.5. User study

We conduct a user study with 30 participants (ages 20s–50s) to further support our evaluation. Participants evaluate two key aspects: prompt fidelity and style fidelity. We selected comparison models based on their quantitative performance: StyleAligned [18] and IP-Adapter [62], which achieved the highest style fidelity, and DB-LoRA [44] and B-LoRA [13], which achieved the highest prompt fidelity. Our method achieves a clearly superior preference

Table 2. Ablation study on Principal Feature Blending (PFB) and Structural Attention Correction (SAC). REP denotes replacement using the style feature F_3^{sty} . The symbol \uparrow indicates that higher is better. The best and second-best results are highlighted in **bold** and underline, respectively.

#	Method	$S_{\text{txt}} \uparrow$	$S_{\text{img}} \uparrow$	$S_{\text{harmonic}} \uparrow$
(a)	Infinity	0.348	0.559	<u>0.429</u>
(b)	Infinity + REP	0.279	0.696	0.398
(c)	Infinity + PFB	0.321	<u>0.631</u>	0.426
(d)	Infinity + PFB + SAC	<u>0.334</u>	0.630	0.437

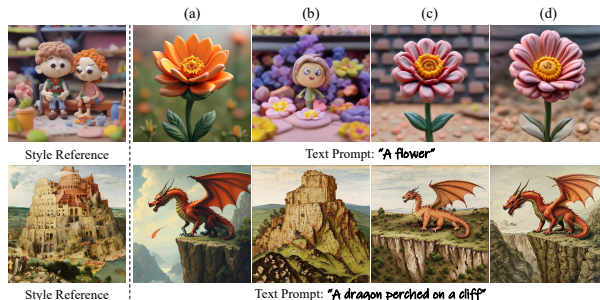


Figure 6. Qualitative ablation study on proposed method. (a)-(d) correspond to the component in Tab. 2.

in prompt fidelity (35.3%) while maintaining competitive style fidelity (32.0%), compared to the other models’ scores of 4.3%, 5.0%, 26.7%, 28.7% (prompt) and 30.7%, 23.3%, 8.3%, 5.7% (style). An example of the interface used in the study is in the Supplementary material.

7. Conclusion

In this work, we introduced a training-free framework for style-personalized image generation that operates on a single reference image and leverages the efficiency of a scale-wise autoregressive model. Through a detailed step-wise analysis of the model’s generation process, we identified a pivotal feature that jointly governs content and style, and further demonstrated—via an SVD-based spectral study—that its dominant singular component captures style-specific variation. Building on these insights, we proposed two lightweight yet effective modules: *Principal Feature Blending*, which provides precise and interpretable style control, and *Structural Attention Correction*, which stabilizes structural consistency during generation. Our method achieves high performance while preserving prompt fidelity, offering a favorable balance. Quantitative and qualitative evaluations confirm that the proposed components operate as intended, enabling faithful style personalization without additional training and with significantly faster inference than existing models.

8. Acknowledgements

This work was supported by LG AI STAR Talent Development Program for Leading Large-Scale Generative AI Models in the Physical AI Domain (RS-2025-25442149), Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (RS-2025-25420118), the InnoCORE program of the Ministry of Science and ICT(N10250156) and the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. RS-2025-02219277, AI Star Fellowship Support Project(DGIST)).

References

- [1] Namhyuk Ahn, Junsoo Lee, Chunggi Lee, Kunhee Kim, Daesik Kim, Seung-Hun Nam, and Kibeom Hong. Dreamstyler: Paint by style inversion with text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 674–681, 2024. 2, 3, 7
- [2] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. Text2live: Text-driven layered image and video editing. In *European conference on computer vision*, pages 707–723. Springer, 2022. 2
- [3] 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. 3
- [4] Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. Coyo-700m: Image-text pair dataset. <https://github.com/kakaobrain/coyo-dataset>, 2022. 2
- [5] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11315–11325, 2022. 3
- [6] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023. 2, 3
- [7] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3558–3568, 2021. 2
- [8] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024. 3
- [9] Jiwoo Chung, Sangeek Hyun, and Jae-Pil Heo. Style injection in diffusion: A training-free approach for adapting large-scale diffusion models for style transfer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8795–8805, 2024. 3
- [10] Yingying Deng, Fan Tang, Weiming Dong, Wen Sun, Feiyue Huang, and Changsheng Xu. Arbitrary style transfer via multi-adaptation network. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2719–2727, 2020. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [12] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 3
- [13] Yarden Frenkel, Yael Vinker, Ariel Shamir, and Daniel Cohen-Or. Implicit style-content separation using b-lora. In *European Conference on Computer Vision*, pages 181–198. Springer, 2024. 2, 3, 7, 8
- [14] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 2
- [15] Jongmin Gim, Jihun Park, Kyounghmin Lee, and Sunghoon Im. Content-adaptive style transfer: A training-free approach with vq autoencoders. In *Proceedings of the Asian Conference on Computer Vision*, pages 2337–2353, 2024. 2
- [16] Jian Han, Jinlai Liu, Yi Jiang, Bin Yan, Yuqi Zhang, Zehuan Yuan, Bingyue Peng, and Xiaobing Liu. Infinity: Scaling bit-wise autoregressive modeling for high-resolution image synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 15733–15744, 2025. 2, 3, 5, 6
- [17] Amir Hertz, Kfir Aberman, and Daniel Cohen-Or. Delta denoising score. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2328–2337, 2023. 3
- [18] Amir Hertz, Andrey Voynov, Shlomi Fruchter, and Daniel Cohen-Or. Style aligned image generation via shared attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4775–4785, 2024. 2, 3, 7, 8
- [19] Kibeom Hong, Seogkyu Jeon, Junsoo Lee, Namhyuk Ahn, Kunhee Kim, Pilhyeon Lee, Daesik Kim, Youngjung Uh, and Hyeran Byun. Aespa-net: Aesthetic pattern-aware style transfer networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22758–22767, 2023. 2
- [20] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. 2
- [21] Jaeseok Jeong, Mingi Kwon, and Youngjung Uh. Training-free content injection using h-space in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5151–5161, 2024. 3
- [22] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *Proceedings of*

- the *IEEE/CVF conference on computer vision and pattern recognition*, pages 10124–10134, 2023. 2
- [23] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vignesh Birodkar, Jimmy Yan, Ming-Chang Chiu, et al. Videopoet: A large language model for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023. 3
- [24] Gihyun Kwon and Jong Chul Ye. Clipstyler: Image style transfer with a single text condition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18062–18071, 2022. 2
- [25] Gihyun Kwon and Jong Chul Ye. Diffusion-based image translation using disentangled style and content representation. *arXiv preprint arXiv:2209.15264*, 2022. 3
- [26] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 2
- [27] Dongxu Li, Junnan Li, and Steven Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Advances in Neural Information Processing Systems*, 36:30146–30166, 2023. 2, 3
- [28] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. *Advances in neural information processing systems*, 30, 2017. 2
- [29] Zhen Li, Mingdeng Cao, Xintao Wang, Zhongang Qi, Ming-Ming Cheng, and Ying Shan. Photomaker: Customizing realistic human photos via stacked id embedding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8640–8650, 2024. 2, 3
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part v 13*, pages 740–755. Springer, 2014. 2
- [31] Chang Liu, Viraj Shah, Aiyu Cui, and Svetlana Lazebnik. Unziplora: Separating content and style from a single image. *arXiv preprint arXiv:2412.04465*, 2024. 2
- [32] Songhua Liu, Tianwei Lin, Dongliang He, Fu Li, Meiling Wang, Xin Li, Zhengxing Sun, Qian Li, and Errui Ding. Adaattn: Revisit attention mechanism in arbitrary neural style transfer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6649–6658, 2021. 2
- [33] Ming Lu, Hao Zhao, Anbang Yao, Yurong Chen, Feng Xu, and Li Zhang. A closed-form solution to universal style transfer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5952–5961, 2019. 2
- [34] Quang-Binh Nguyen, Minh Luu, Quang Nguyen, Anh Tran, and Khoi Nguyen. Csd-var: Content-style decomposition in visual autoregressive models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17013–17023, 2025. 4
- [35] Dae Young Park and Kwang Hee Lee. Arbitrary style transfer with style-attentional networks. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5880–5888, 2019. 2
- [36] Jihun Park, Jongmin Gim, Kyoungmin Lee, Seunghun Lee, and Sunghoon Im. Style-editor: Text-driven object-centric style editing. *arXiv preprint arXiv:2408.08461*, 2025. 2
- [37] Jihun Park, Jongmin Gim, Kyoungmin Lee, Minseok Oh, Minwoo Choi, Jaeyul Kim, Woo Chool Park, and Sunghoon Im. A training-free style-aligned image generation with scale-wise autoregressive model. *arXiv preprint arXiv:2504.06144*, 2025. 3, 4
- [38] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 2
- [39] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [40] 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. 2, 3
- [41] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021. 2
- [42] 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. 2
- [43] 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
- [44] Simo Ryu. Low-rank adaptation for fast text-to-image diffusion fine-tuning, 2022. URL <https://github.com/cloneofsimo/lorax>, 10:19, 2022. 2, 3, 7, 8
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 2
- [46] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in neural information processing systems*, 35:25278–25294, 2022. 2
- [47] Viraj Shah, Nataniel Ruiz, Forrester Cole, Erika Lu, Svetlana Lazebnik, Yuanzhen Li, and Varun Jampani. Ziplora: Any subject in any style by effectively merging loras. In *European Conference on Computer Vision*, pages 422–438. Springer, 2024. 2, 3

- [48] Joonghyuk Shin, Alchan Hwang, Yujin Kim, Daneul Kim, and Jaesik Park. Exploring Multimodal Diffusion Transformers for Enhanced Prompt-based Image Editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 6
- [49] 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, 3, 7
- [50] Haotian Tang, Yecheng Wu, Shang Yang, Enze Xie, Junsong Chen, Junyu Chen, Zhuoyang Zhang, Han Cai, Yao Lu, and Song Han. Hart: Efficient visual generation with hybrid autoregressive transformer. *arXiv preprint arXiv:2410.10812*, 2024. 2, 3
- [51] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024. 3
- [52] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. 3
- [53] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1921–1930, 2023. 6
- [54] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 3
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [56] Anton Voronov, Denis Kuznedelev, Mikhail Khoroshikh, Valentin Khrulkov, and Dmitry Baranchuk. Switti: Designing scale-wise transformers for text-to-image synthesis, 2025. 3
- [57] Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15943–15953, 2023. 2, 3
- [58] Yi Wu, Lingting Zhu, Shengju Qian, Lei Liu, Wandu Qiao, Lequan Yu, and Bin Li. Stylear: Customizing multimodal autoregressive model for style-aligned text-to-image generation. *arXiv preprint arXiv:2505.19874*, 2025. 7
- [59] Peng Xing, Haofan Wang, Yanpeng Sun, Qixun Wang, Xu Bai, Hao Ai, Renyuan Huang, and Zechao Li. Csgo: Content-style composition in text-to-image generation. *arXiv preprint arXiv:2408.16766*, 2024. 7
- [60] Serin Yang, Hyunmin Hwang, and Jong Chul Ye. Zero-shot contrastive loss for text-guided diffusion image style transfer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22873–22882, 2023. 3
- [61] Yuan Yao, Jianqiang Ren, Xuansong Xie, Weidong Liu, Yong-Jin Liu, and Jun Wang. Attention-aware multi-stroke style transfer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1467–1475, 2019. 2
- [62] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 7, 8
- [63] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gungjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022. 6
- [64] Gong Zhang, Kihyuk Sohn, Meera Hahn, Humphrey Shi, and Irfan Essa. Finestyle: Fine-grained controllable style personalization for text-to-image models. *Advances in Neural Information Processing Systems*, 37:52937–52961, 2024. 2, 3, 6, 7