

# Z-SASLM: Zero-Shot Style-Aligned SLI Blending Latent Manipulation

Alessio Borgi\* Luca Maiano Irene Amerini

Sapienza University of Rome

Department of Computer, Control and Management Engineering

{borgi.1952442@studenti, maiano@diag, amerini@diag}.uniroma1.it

## Abstract

We introduce Z-SASLM, a Zero-Shot Style-Aligned SLI (Spherical Linear Interpolation) Blending Latent Manipulation pipeline that overcomes the limitations of current multi-style blending methods. Conventional approaches rely on linear blending, assuming a flat latent space leading to suboptimal results when integrating multiple reference styles. In contrast, our framework leverages the non-linear geometry of the latent space by using SLI Blending to combine weighted style representations. By interpolating along the geodesic on the hypersphere, Z-SASLM preserves the intrinsic structure of the latent space, ensuring high-fidelity and coherent blending of diverse styles—all without the need for fine-tuning. We further propose a new metric, Weighted Multi-Style DINO VIT-B/8, designed to quantitatively evaluate the consistency of the blended styles. While our primary focus is on the theoretical and practical advantages of SLI Blending for style manipulation, we also demonstrate its effectiveness in a multi-modal content fusion setting through comprehensive experimental studies. Experimental results show that Z-SASLM achieves enhanced and robust style alignment. The implementation code can be found at: <https://alessioborgi.github.io/Z-SASLM.github.io/>.

## 1. Introduction

Text-to-image generation has advanced rapidly, moving from GANs[1] and its numerous variants[2–7] to powerful diffusion-based models [8, 9]. The rapid development of large-scale text-to-image (T2I) models such as DALL-E [10], MidJourney [11], Stable Diffusion [12] and others [13–15] has revolutionized creative industries by enabling the generation of high-quality, diverse visual outputs from textual descriptions. Despite these advances, achieving consistent *Style Alignment*—the ability to maintain a coherent visual style across multiple generated images—remains a

\*Corresponding author: borgi.1952442@studenti.uniroma1.com




Figure 1. Medieval-Cubism SLI Blending (2-styles)

significant challenge. For example, a designer aiming to blend the distinct aesthetics of cubism and baroque across visuals must often resort to fine-tuning models on specific styles. This approach limits the blending of the styles present in the fine-tuning dataset and confines current methods to single-style references.

In this work, we propose a novel Z-SASLM architecture based on *Spherical Linear Interpolation (SLI) Blending* approach for multi-reference style conditioning, a technique

able to interpolate along the geodesic on the hypersphere, preserving the image manifold’s geometric properties and ensuring smooth, coherent blending between styles. Importantly, our method eliminates the need for fine-tuning, enabling zero-shot style alignment directly during generation.

**Paper Contributions.** Our main contributions are the following:

- **Z-SASLM architecture based on SLI Blending for Multi-Reference Style Conditioning:** We introduce a SLI-based blending technique that leverages the latent space’s non-linear geometry to combine multiple reference-style images in a weighted manner. This approach overcomes the limitations of linear blending and bypasses the need for fine-tuning.
- **Weighted Multi-Style DINO ViT-B/8 Weighted Metric:** We propose a new metric designed to evaluate the consistency of the style in a set of generated images, effectively quantifying the contributions of multiple blended styles.
- **Multi-Modal Content Fusion Ablation:** We conduct comprehensive ablation studies using multi-modal content fusion—integrating image, audio, and other modalities—to validate the improvements in style alignment even in a multi-modal setting.

## 2. Related Work

**Style Alignment in Text-to-Image Models.** Despite significant advances by models such as DALL·E [10], Stable Diffusion [12], and Imagen [8], achieving consistent style alignment remains a challenging task. Most state-of-the-art T2I models are designed for single-image generation and optimized for a single style reference, often resulting in inconsistencies when generating a series of images. Early approaches to style alignment can be broadly divided into *Fine-Tuning-Based* and *Latent Space Manipulation* methods. Fine-tuning techniques, like StyleDrop [16] and DreamBooth [17], require adapting the model to specific styles, which is computationally intensive, restricts the range of achievable styles, and limits scalability. Latent space manipulation methods, exemplified by StyleGAN [18] and Style-FiT [19], enable zero-shot style transfer but are predominantly designed for single style adaptation and often struggle with content-style disentanglement when blending multiple styles. In essence, while these methods generate individual images conditioned on a style reference, they do not enforce uniform style consistency across a set of generated images—as if they were all created by the same artist. Approaches such as StyleAligned[20] is the most similar approach to our idea, using shared attention mechanisms to enforce style consistency in a zero-shot manner; however, they remain confined to a single style reference and do not support weighted blending of multiple styles.

**Multi-Reference Style Blending.** Traditional methods to blend multiple reference styles, such as those employed in StyleGAN2-ADA [21], permit style mixing by directly manipulating the latent space. This primarily relies on simple linear combinations designed for GAN-based architectures. This approach suffers from several limitations: it assumes a flat (Euclidean) latent space that fails to capture the curved geometry of high-dimensional representations, often resulting in abrupt transitions, artifacts, and incoherent style mixtures. In contrast, our work introduces a multi-reference-weighted style blend framework tailored to diffusion-based models, with at its core, SLI Blending, which interpolates along the geodesic on the hypersphere. This approach preserves the intrinsic structure of the latent manifold, ensuring that the weighted combination of style representations remains within a semantically meaningful region. By overcoming the inherent limitations of conventional, linear-based methods, our SLI Blending technique achieves a high-fidelity, coherent fusion of diverse styles, enabling smooth and consistent multi-style integration.

**Multi-Style Consistency Evaluation.** Evaluating the consistency of blended styles poses unique challenges. While metrics such as DINO ViT [22] have been successfully employed to assess style consistency, they are typically tailored to single-style scenarios and do not account for the nuances of weighted multi-style blending. To this end, we propose the Weighted Multi-Style DINO ViT-B/8 metric, an enhanced evaluation tool that extends traditional DINO ViT metrics to effectively quantify the coherence and contribution of multiple blended styles.

## 3. Method Overview

The Z-SASLM architecture, illustrated in Fig. 2, consists of three primary modules—Reference Images Encoding & Blending, Text Encoding, and the StyleAligned Image Generation process—augmented by an optional Multi-Modal Content Fusion module. When this optional module is not employed, a simple caption (denoted as *Single-Content Textual prompt*) indicating the desired scene we want to generate can serve as a substitute. The Multi-Modal Content Fusion module, on the other hand, provides a *Multi-Content Textual Prompt* that can aggregate diverse inputs such as images, audio, music, and weather data. It utilizes T5-based rephrasing to merge these modalities into a unified textual prompt, which is then forwarded to the Text Encoding module. Here, the prompt (either coming from the Multi-Modal Content Fusion module or from a simple caption) is tokenized and encoded using CLIP, producing embeddings for both positive and negative prompts. In parallel, the Reference Image Encoding & Blending module extracts latent vectors from multiple reference styles via a VAE encoder. These latent vectors are blended using our proposed SLI (Spherical Linear Interpolation) Blending approach. Fi-

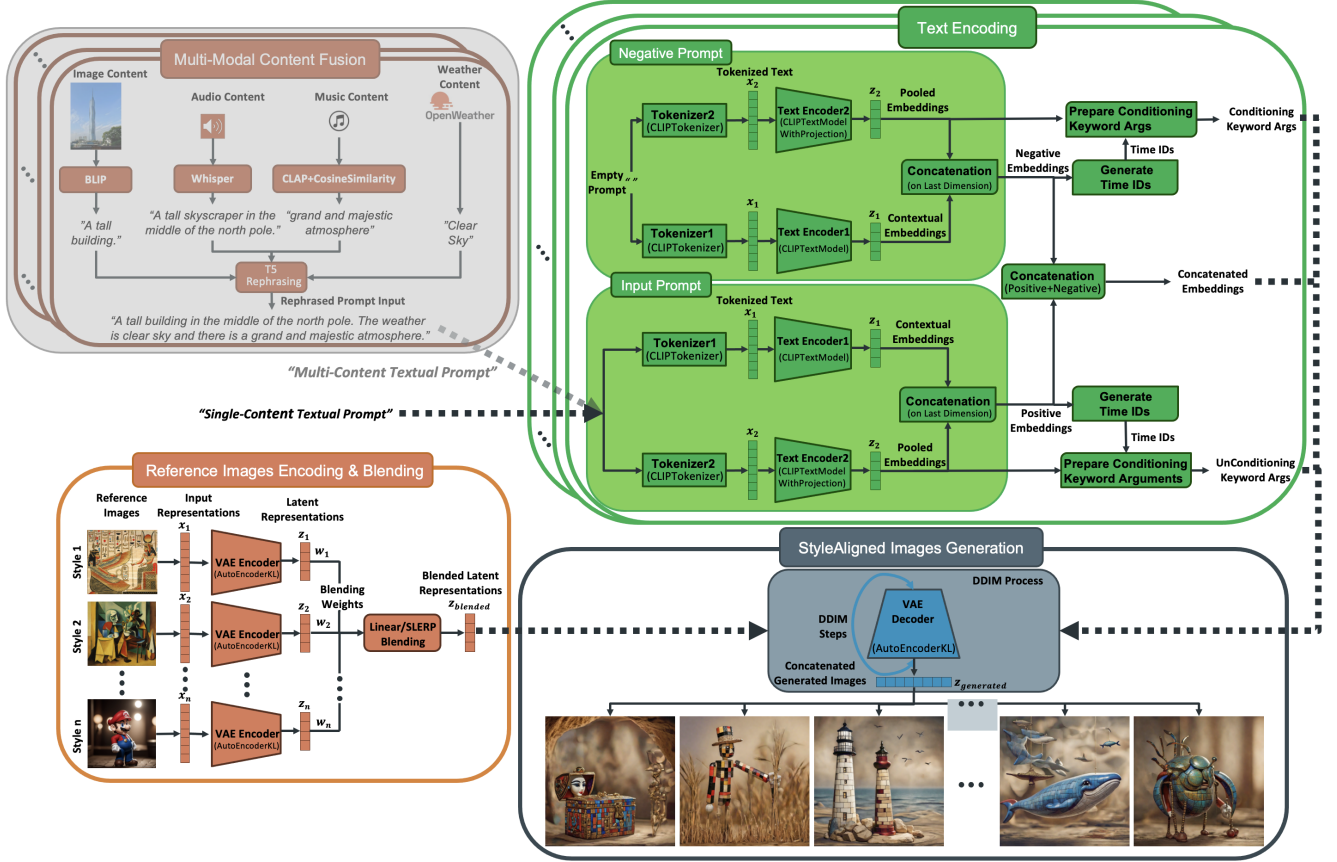


Figure 2. Overview of the Z-SASLM Architecture.

nally, the unified style representation is combined with the textual embeddings and provided to the StyleAligned image generation process, resulting in images that exhibit coherent and consistent style alignment.

### 3.1. Background

To achieve style alignment during generation without fine-tuning, we adopt key components from StyleAligned [20]: Adaptive Instance Normalization (AdaIN) and Shared Attention.

**Adaptive Instance Normalization (AdaIN).** AdaIN [23] aligns the feature statistics of a generated image with those of a reference style image. Let  $g = f(\mathbf{g})$  and  $s = f(\mathbf{s})$  be the feature maps for the generated image and reference style, respectively, and define  $\mu(\cdot)$  and  $\sigma(\cdot)$  as the mean and standard deviation. Then AdaIN is:

$$\text{AdaIN}(g, s) = \sigma(s) \left( \frac{g - \mu(g)}{\sigma(g)} \right) + \mu(s) \quad (1)$$

**Shared Attention.** Shared Attention propagates style information across multiple images by computing pairwise similarities. For a set  $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ , with query  $Q(\cdot)$ , key  $K(\cdot)$ ,

and value  $V(\cdot)$  features, the attention weights are:

$$A_{ij} = \frac{\exp(Q(f(\mathbf{x}_i)) \cdot K(f(\mathbf{x}_j))^T / \sqrt{d})}{\sum_{k=1}^n \exp(Q(f(\mathbf{x}_i)) \cdot K(f(\mathbf{x}_k))^T / \sqrt{d})} \quad (2)$$

which update the style representation as:

$$f(\mathbf{x}_i)' = \sum_{j=1}^n A_{ij} V(f(\mathbf{x}_j)) \quad (3)$$

In summary, queries and keys are first normalized via AdaIN using the reference style’s statistics:

$$\hat{Q}_i = \text{AdaIN}(Q_i, Q_{\text{ref}}), \quad \hat{K}_i = \text{AdaIN}(K_i, K_{\text{ref}}) \quad (4)$$

The final style-aligned attention map is then computed by concatenating keys and values from both the reference image and the current image:

$$\text{Attention}(\hat{Q}_i, [K_{\text{ref}}, \hat{K}_i], [V_{\text{ref}}, V_i]) \quad (5)$$

By incorporating these mechanisms into our architecture, we enable zero-shot style alignment during generation, eliminating the need for fine-tuning.



### 3.2. Multi-Reference Weighted Style Blending

Inspired by the style-mixing approach of StyleGAN2-ADA [21], which linearly interpolates latent codes for GAN-based models, we adapt a similar idea to diffusion-based models. However, because diffusion models typically operate in a non-Euclidean latent space, naively applying linear interpolation can lead to suboptimal or inconsistent style transitions. In this section, we first present linear blending as a baseline, then introduce our Spherical Linear Interpolation (SLI) Blending to address the limitations of linear mixing.

#### 3.2.1. Linear Weighted Style Blending

Suppose we have  $k$  reference style images  $\{s_1, s_2, \dots, s_k\}$  mapped to latent vectors  $\{z_1, z_2, \dots, z_k\}$ , each weighted by  $\{w_1, w_2, \dots, w_k\}$  with  $\sum_{i=1}^k w_i = 1$ . A straightforward approach, similar to StyleGAN2-ADA [21], is to form a *linear* combination of these style vectors:

$$z_{\text{blend}}^{\text{linear}} = \sum_{i=1}^k w_i z_i \quad (6)$$

We then use  $z_{\text{blend}}$  as a conditioning latent in the diffusion process. Although simple and intuitive, linear interpolation assumes a flat (Euclidean) geometry; in the curved, high-dimensional latent spaces of diffusion models, it can yield abrupt transitions or washed-out stylistic details. Consequently, by making several experiments, we often observe noticeable artifacts in the generated images, as the number of reference images augments, as illustrated in Fig. 3. These artifacts arise because the linear path between style vectors frequently falls off the meaningful regions of the latent manifold, causing a loss of fine-grained stylistic cues and reducing overall fidelity.

#### 3.2.2. SLI Blending

We propose using *Spherical Linear Interpolation (SLI) Blending* to better preserve style fidelity. The idea of spherical interpolation was first introduced by Shoemake *et al.* [24] and consists in following the geodesic on a hypersphere rather than a straight line in Euclidean space, as illustrated in Fig. 4. By respecting the manifold’s curvature, SLI produces smoother and more coherent blends of multiple styles compared to linear interpolation. Linear blending computes the weighted sum which, even if  $z_1$  and  $z_2$  are unit vectors, does not guarantee that  $z_{\text{blend}}^{\text{linear}}$  will also be of unit length. This deviation from the hypersphere can push the result away from high-density regions of the latent space, resulting in abrupt transitions and loss of fine stylistic details. Mathematically, for two unit vectors with angular separation  $\omega$ , their Euclidean (chord) distance is:

$$d_{\text{linear}} = 2 \sin\left(\frac{\omega}{2}\right) \quad (7)$$

while the true geodesic (arc) distance on the hypersphere is:

$$d_{\text{geodesic}} = d_{\text{sl}} = \omega \quad (8)$$

Since  $2 \sin\left(\frac{\omega}{2}\right) < \omega$  for  $\omega > 0$ , linear interpolation underestimates the true separation between styles. In contrast, SLI precisely follows the arc, ensuring that the interpolation accurately reflects the perceptual distance between the styles.

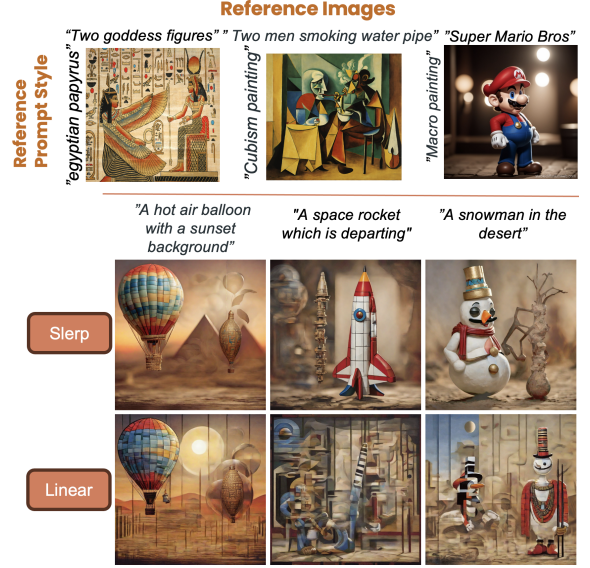


Figure 3. Linear vs. SLI Blending with three styles. Linear interpolation can fall off the meaningful latent manifold, introducing artifacts. On the other hand, SLERP always provides robust results.

**Definition.** Starting with the simplest case in which we have *two style vectors*  $z_1$  and  $z_2$  with weights  $w_1$  and  $w_2$ , let

$$t = \frac{w_2}{w_1 + w_2}, \quad \omega = \arccos\left(\frac{z_1 \cdot z_2}{\|z_1\| \|z_2\|}\right) \quad (9)$$

Assuming that the latent vectors are normalized ( $\|z_1\| = \|z_2\| = 1$ ), SLI is defined as:

$$\text{SLI}(t, z_1, z_2) = \frac{\sin((1-t)\omega)}{\sin(\omega)} z_1 + \frac{\sin(t\omega)}{\sin(\omega)} z_2 \quad (10)$$

When  $w_1 = w_2$ ,  $t = 0.5$  yields an equal contribution; otherwise, the interpolation skews toward the style with the higher weight.

For  $k > 2$  styles, we can simply extend SLI blending by iteratively combining style vectors. However, because SLI is a nonlinear operation defined on the hypersphere, its iterative application is not associative; that is, the final blended vector can depend on the order in which the style vectors are



Table 1. DDIMScheduler Configuration

Parameter	Value
beta_start	0.00085
beta_end	0.012
beta_schedule	scaled_linear
clip_sample	False
set_alpha_to_one	False




Figure 4. 2-style Linear vs. SLI in the latent space Illustration. The chord (in red) represents naive linear blending between two reference vectors  $\mathbf{z}_1$  and  $\mathbf{z}_2$ , yielding  $\mathbf{z}_{\text{blend}}^{\text{lin}}$ . In contrast, Spherical Linear Interpolation (in green) follows the geodesic on the unit sphere, producing  $\mathbf{z}_{\text{blend}}^{\text{SLI}}$ . The angle  $\omega$  denotes the separation between  $\mathbf{z}_1$  and  $\mathbf{z}_2$ , with  $d_{\text{linear}} < d_{\text{sli}}$  on the unit circle.

combined. To address this, we adopt a semantically meaningful ordering strategy by *sorting the reference styles in descending order of their weights*. This way, the most dominant styles are incorporated first, ensuring that their influence is robustly preserved in the cumulative blend, providing a stable, reproducible blending sequence and mitigating potential artifacts due to arbitrary ordering. More formally, let  $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k$  be the latent style vectors sorted such that  $w_1 \geq w_2 \geq \dots \geq w_k$ . If  $\mathbf{z}_{1\dots i}$  denotes the cumulative blend of the first  $i$  styles, we incorporate the  $(i+1)$ -th style with weight  $w_{i+1}$  as follows:

$$\mathbf{z}_{1\dots(i+1)} = \text{SLI}\left(\frac{w_{i+1}}{\sum_{m=1}^{i+1} w_m}, \mathbf{z}_{1\dots i}, \mathbf{z}_{i+1}\right) \quad (11)$$

Repeating this process yields the final blended latent vector  $\mathbf{z}_{\text{blend}}^{\text{SLI}}$ .

### 3.2.3. Dynamic Style-Aligned Arguments Scaling

Even with SLI, certain styles can dominate the blending process due to their more pronounced activations in the latent space. We refer to styles that tend to generate such strong responses as “famous,” while those with standard activations are considered “normal.” In our approach, we discriminate between these two categories by examining the norm of the latent key  $\mathbf{k}$  for each style. Specifically, if  $\|\mathbf{k}\|$  exceeds a predefined threshold  $T = 0.5$ , the style is

Table 2. SDXL Pipeline Configuration

Parameter	Value
torch_dtype	torch.float16
variant	fp16
use_safetensors	True
scheduler	DDIMScheduler
device	cuda
num_inference_steps	50

classified as famous. This criterion is well-justified because the dot product in the attention mechanism scales with the norms:

$$\langle Q, K \rangle \propto \|Q\| \|K\| \quad (12)$$

so a higher  $\|\mathbf{k}\|$  directly results in disproportionately high attention scores, causing famous styles to overshadow others. This issue was also observed in the StyleAligned paper [20]. To mitigate this imbalance, we rescale the attention scores for each style reference by applying a style-dependent normalization. Specifically, we adjust the original attention score  $A_{\text{original}}$  as:

$$A_{\text{normalized}} = A_{\text{original}} \times \sigma + \mu \quad (13)$$

where the shift  $\mu$  and scale  $\sigma$  are determined by the style’s classification (as proposed in [20]):

$$\{\mu, \sigma\} = \begin{cases} \{\log(2), 1\} & \text{style} == \text{normal ("n")} \\ \{\log(1), 0.5\} & \text{style} == \text{famous ("f")} \end{cases} \quad (14)$$

This normalization dampens the excessive influence of famous styles by reducing their scale while preserving or even slightly boosting the contribution of normal styles through an appropriate shift.

### 3.3. Weighted Multi-Style DINO VIT-B/8

Evaluating style consistency in text-to-image (T2I) models requires a robust metric that captures both fine-grained visual details and higher-level semantic style characteristics. The DINO (Distillation with No Labels) VIT-B/8 model [22]—a self-supervised vision transformer—has proven effective in this regard. Thanks to its multi-headed self-attention layers and an 8-patch input scheme, DINO VIT-B/8 extracts feature representations that encapsulate subtle stylistic nuances without requiring manual labels.

To assess the alignment between a generated image and a reference style, we first compute the cosine similarity between their corresponding feature embeddings. Let  $\mathbf{z}_{\text{gen}}$  denote the embedding of the generated image and  $\mathbf{z}_{\text{ref}}$  that of a reference style image. The cosine similarity is defined as:

$$\text{CS}(\mathbf{z}_{\text{gen}}, \mathbf{z}_{\text{ref}}) = \frac{\mathbf{z}_{\text{gen}} \cdot \mathbf{z}_{\text{ref}}}{\|\mathbf{z}_{\text{gen}}\| \|\mathbf{z}_{\text{ref}}\|} \quad (15)$$

While DINO VIT-B/8 excels at comparing a generated image against a single style reference, it does not directly

support the evaluation of multi-style consistency. To address this, we extend the conventional cosine similarity into a weighted multi-style metric. Given the feature embeddings of  $k$  reference style images and a generated image, we define the weighted multi-style similarity score as:

$$S_{\text{multi-style}}(\mathbf{z}_{\text{gen}}, \mathbf{z}_1, \dots, \mathbf{z}_k) = \sum_{i=1}^k w_i \cdot \text{CS}(\mathbf{z}_{\text{gen}}, \mathbf{z}_i) \quad (16)$$

This formulation effectively combines the individual style alignments into a single metric that reflects the overall style consistency of the generated image with respect to multiple references. More specifically, if we denote the cosine similarity for each reference as  $s_i = \text{CS}(\mathbf{z}_{\text{gen}}, \mathbf{z}_i)$ , then the final metric can be expressed as a weighted average:

$$S_{\text{multi-style}} = \sum_{i=1}^k w_i s_i \quad (17)$$

## 4. Experiments

We conduct our experiments on Stable Diffusion XL (SDXL), specifically the pre-trained model ‘*stabilityai/stable-diffusion-xl-base-1.0*’ [25], following the setup of Hertz *et al.* [20]. However, our approach is model-agnostic and can be adapted to other diffusion-based or generative models. We use a DDIM (Denoising Diffusion Implicit Models) [26] scheduler (configuration in Tab. 1) and initialize the SDXL pipeline with parameters in Tab. 2. All experiments are performed on GPUs in mixed-precision mode for efficiency.

### 4.1. Weighted Multi-Style DINO VIT-B/8: Linear vs SLI Interpolation

We first compare our SLI Blending against the simpler Linear Blending approach, which is inspired by the linear interpolation used in StyleGAN2-ADA [21] but extended to diffusion models. To evaluate style consistency, we use the proposed Weighted Multi-Style DINO VIT-B/8 metric (indicated as  $\text{WMS}_{\text{DINO-VIT-B/8}}$  in the table) and also report the CLIP Score [27] for image-text alignment. In addition, we compute and report the mean similarity scores for the Medieval and Cubism reference images, denoted as  $\text{MS}_{\text{med}}$  and  $\text{MS}_{\text{cub}}$ , respectively. These metrics quantify the alignment of generated images with each reference style individually, providing insight into how well each style is preserved in the blended output.

**Setup.** We generate images by blending two reference styles (e.g., *Medieval* and *Cubism*) using several weight configurations: (0, 1), (0.15, 0.85), (0.25, 0.75), (0.5, 0.5),

(0.75, 0.25), (0.85, 0.15), and (1, 0), and we fix the guidance scale to 20. In our experiments, the weight configuration, e.g.  $(w_{\text{med}}, w_{\text{cub}})$ , determines the relative influence of the Medieval and Cubism styles on the generated image. For example, (0, 1) yields full Cubism and (1, 0) full Medieval; (0.5, 0.5) produces an even blend, while (0.25, 0.75) or (0.75, 0.25) skew the output toward Cubism or Medieval, respectively. Inspired by the StyleAligned[20] approach, we also created a dataset for the textual input prompts, composed of seven sets of nine images, using ChatGPT. This dataset comprises a diverse collection of stylistic descriptions and artistic references, enabling systematic evaluation of our blending methods across a wide range of style fusion scenarios.

**Results.** Table 3 summarizes our findings. The rows {0, 1} and {1, 0} represent the *StyleAligned* [20] baseline with no multi-style blending. From Tab. 3, SLI outperforms Linear interpolation in multi-style alignment ( $\text{WMS}_{\text{DINO-VIT-B/8}}$ ) while maintaining better CLIP scores. The differences become more pronounced in blended scenarios (e.g., {0.5, 0.5}) where linear blending struggles to stay on the latent manifold, leading to suboptimal style fusion.

### 4.2. Multi-Modal Content Fusion Module Inclusion

Although our primary focus is on multi-style blending, we also evaluate an optional *Multi-Modal Content Fusion* module (see Fig. 2) to demonstrate the advantages of enriching the input prompt beyond simple text, defined in our architecture as “*Single-Content Textual Prompt*”. In many creative tasks, a single-context prompt may lack the nuance necessary to capture the full spectrum of an intended artistic vision. By incorporating additional modalities—such as visual cues, audio transcripts, musical mood descriptors, and real-time weather data—we provide a richer and more informative conditioning signal that enhances style alignment and improves image quality.

In our approach, each modality is first converted into text: a photo is described via *BLIP*[28]-based *Image-to-Text*, spoken content is transcribed with *Whisper*[29], music is interpreted via *CLAP*[30] combined with *cosine similarity* matching, and weather data is retrieved through the *OpenWeather API*[31]. This process has been employed to continue avoiding to perform the fine-tuning step. These diverse textual snippets are then concatenated and further condensed using *T5*[32]-based *paraphrasing* to produce a compact, yet semantically rich, “*Multi-Content Textual Prompt*” that is fed into the SDXL pipeline.

As shown in Fig. 5, our Z-SASLM Blending method maintains coherent style alignment even under these richer conditions, demonstrating the advantage of multi-modal fusion in guiding the generation process.

Style Weights	Linear (StyleGAN2-ADA[21]-adapted)				Z-SASLM (Ours)			
$\{w_{\text{med}}, w_{\text{cub}}\}$	$\text{MS}_{\text{med}}$	$\text{MS}_{\text{cub}}$	$\text{WMS}_{\text{DINO-VIT-B/8}}$	$\text{CLIP}_{\text{score}}$	$\text{MS}_{\text{med}}$	$\text{MS}_{\text{cub}}$	$\text{WMS}_{\text{DINO-VIT-B/8}}$	$\text{CLIP}_{\text{score}}$
$\{0, 1\}^*$	-	0.47552	0.47552	0.30280	-	0.47552	0.47552	0.30280
$\{0.15, 0.85\}$	0.32466	0.42683	0.41151	<b>0.31534</b>	0.32595	0.47072	<b>0.44900</b>	0.31049
$\{0.25, 0.75\}$	0.35550	0.42250	0.40575	0.31420	0.33046	0.45447	<b>0.42347</b>	<b>0.31657</b>
$\{0.5, 0.5\}$	0.34905	0.37881	0.36393	0.29232	0.36150	0.42156	<b>0.39153</b>	<b>0.31434</b>
$\{0.75, 0.25\}$	0.35798	0.38327	<b>0.36430</b>	0.31752	0.34648	0.35099	0.34760	<b>0.31911</b>
$\{0.85, 0.15\}$	0.35513	0.40860	0.36315	<b>0.32381</b>	0.36513	0.38286	<b>0.36779</b>	0.31499
$\{1, 0\}^*$	0.29891	-	0.29891	0.30570	0.29891	-	0.29891	0.30570

Table 3. Weighted Multi-Style DINO VIT-B/8: Linear vs. Z-SASLM. \*: Rows  $\{0, 1\}$  and  $\{1, 0\}$  indicate no blending (i.e., StyleAligned [20]). We omit from the comparison the basic SDXL results as we are comparing only style-alignment results.

Table 4. Ablation Study: Attention Scaling vs. Non-Scaling (Z-SASLM’s SLI  $\{0.5, 0.5\}$  Blending)

Scaling Parameters ( <i>Cubism</i> )	$\text{MS}_{\text{med}}$	$\text{MS}_{\text{cub}}$	$ \text{MS}_{\text{med}} - \text{MS}_{\text{cub}} $	$\text{WMS}_{\text{DINO-VIT-B/8}}$	$\text{CLIP}_{\text{score}}$
$\{\log(1), 0.125\}$	0.36096	0.40155	0.04059	0.38126	0.31817
$\{\log(1), 0.25\}$	0.35660	0.39909	0.04249	0.37784	0.28070
$\{\log(1), 0.5\}$	0.36150	0.42156	0.06006	<b>0.39153</b>	<b>0.31434</b>
$\{\log(1), 0.75\}$	0.36272	0.41224	0.04952	0.38748	0.31706
$\{\log(2), 1\}^\dagger$	0.33432	0.45480	0.12048	0.39456	0.30885

$^\dagger$ : Without Attention Scaling

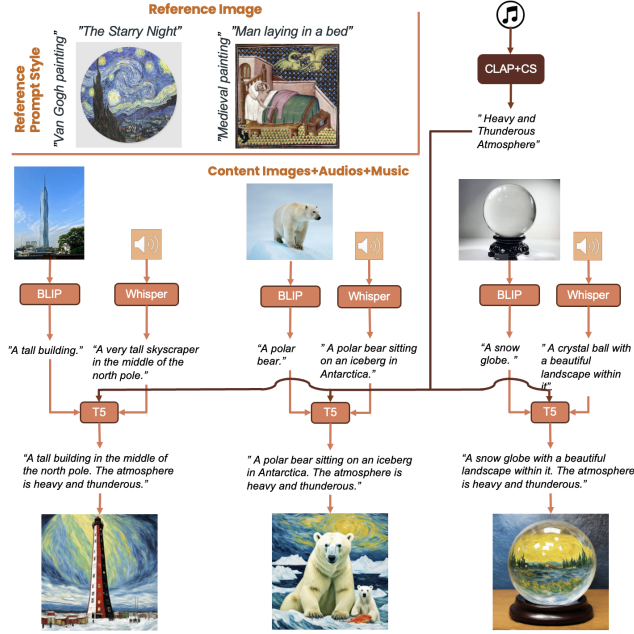


Figure 5. Z-SASLM with Multi-Modal Content Fusion. Our SLI blending (with  $(w_{\text{med}}, w_{\text{vangogh}}) = (0.15, 0.85)$ ) preserves style alignment despite varied contextual cues (image, audio, music). As you can observe from the photos, the impact of the other modalities on the generated image is tangible (e.g., the background of the images is “heavy and thunderous” as the music modalities suggest). In the meanwhile, information given through Audio further allows to be more detailed in the generation.

### 4.3. Scaling vs Non-Scaling

We next investigate the impact of attention scaling ( $\mu, \sigma$ ) on Z-SASLM’s SLI blending. Using an equal weight configuration  $\{0.5, 0.5\}$  for two reference styles, we vary the scaling parameters for the more prominent style (*Cubism*). Table 4 shows that without scaling, the difference between mean similarities  $|\text{MS}_{\text{med}} - \text{MS}_{\text{cub}}|$  is significantly larger, implying style imbalance. Rescaling helps mitigate dominance by famous styles, ensuring a more balanced blend. CLIP scores remain fairly stable. You can also observe this behavior from Figure 6, noting the predominance of the Cubism style w.r.t. the medieval one during image generation.

### 4.4. Guidance Ablation

Finally, we vary the guidance scale (5 to 30) under SLI blending with weights  $\{0.5, 0.5\}$ . Table 5 indicates that higher guidance often increases the Weighted Multi-Style DINO VIT-B/8 metric but may also increase  $|\text{MS}_{\text{med}} - \text{MS}_{\text{cub}}|$ . This suggests a trade-off: stronger guidance can boost style fidelity but might amplify one style more than the other. Overall, moderate guidance (15–20) strikes a good balance. This behavior can also be observed directly through Figure 7, where you can observe how increasing the guidance scale progressively amplifies the stylistic influence from the two reference images. At low guidance (5), each generated image remains relatively faithful to its textual prompt. As guidance rises (10–15), characteristic shapes, colors, and patterns drawn from the reference images become more evident. By higher guidance levels (20–25), the stylization is more aggressive; the



Z-SASLM {0.5, 0.5}					
Guidance	MS <sub>med</sub>	MS <sub>cub</sub>	MS <sub>med</sub> − MS <sub>cub</sub>	WMS <sub>DINO-ViT-B/8</sub>	CLIP <sub>score</sub>
5	0.37721	0.38623	0.00902	0.38172	0.31420
10	0.36149	0.42156	0.06007	0.39153	0.31434
15	0.35610	0.43315	0.07705	0.39463	0.31681
20	0.38218	0.45844	0.07626	<b>0.42031</b>	<b>0.31656</b>
25	0.36966	0.44971	0.08005	0.40968	0.31554
30	0.36350	0.46818	0.10468	0.41584	0.31316

Table 5. Guidance Ablation Study for Z-SASLM {0.5, 0.5} Blending.

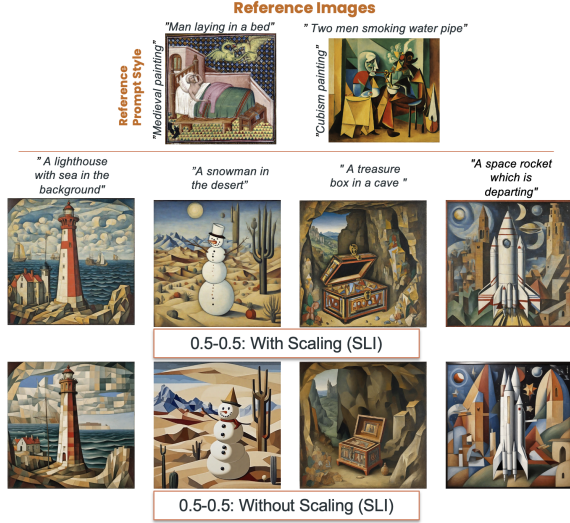


Figure 6. Attention score rescaling effect under SLI blending.

original content is still recognizable, but geometric forms, bold outlines, and subdued palettes—drawn from the references—dominate the overall look. In other words, strong guidance enhances style fidelity but can overshadow fine details of the prompt, creating images that lean more heavily toward the shared artistic signature of the references.

## 5. Conclusions

We have presented the novel Z-SASLM framework for multi-style blending in diffusion models, focusing on a Spherical Linear Interpolation (SLI) approach that respects the non-Euclidean geometry of the latent space. Our experiments show that Z-SASLM outperforms the Linear baseline (inspired by StyleGAN2-ADA) across various style weight configurations and under multi-modal content prompts. Additionally, we introduced a Weighted Multi-Style DINO VIT-B/8 metric to quantify style consistency, demonstrating the superiority of SLI in navigating complex latent manifolds. Although our main emphasis is on multi-style blending, we also illustrated that Z-SASLM’s solution remains robust when additional modalities (audio, music, weather)

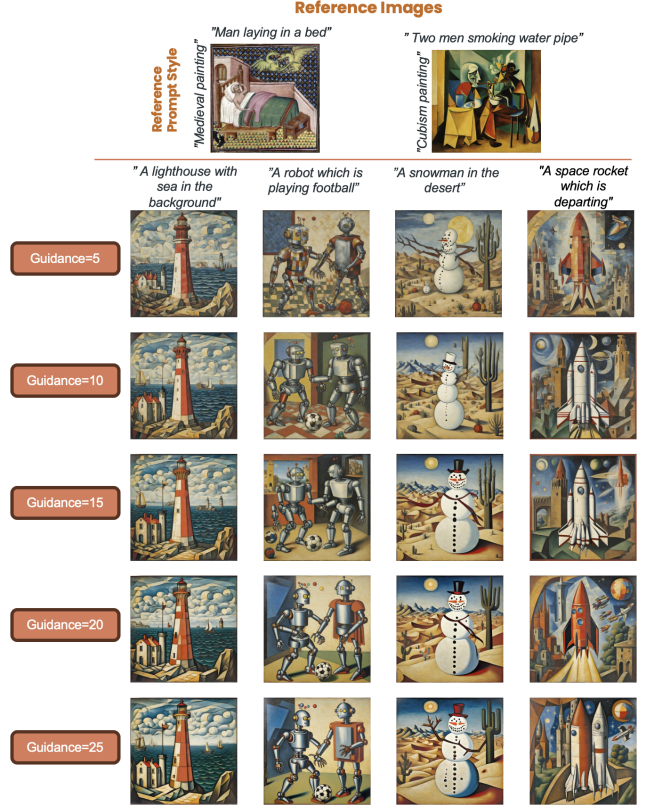


Figure 7. Z-SASLM: Guidance Ablation showing the effect of different guidance scaling factors.

are fused into the prompt. Future work will explore more efficient attention mechanisms and extended style references for large-scale applications.

## 6. Acknowledgement

This study has been partially supported by SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU.

## References

- [1] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. 1
- [2] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks, 2017. 1
- [3] Jiahui Yu, Yuanzhong Xu, Chitwan Saharia, William Chan, Han Zhang, Mohammad Norouzi, Jason Baldridge, and Ting Chen. Parti: Pathways autoregressive text-to-image model, 2022. Parti.
- [4] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks, 2017.
- [5] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaoogang Wang, Xiaolei Huang, and Dimitris Metaxas. StackGAN++: Realistic image synthesis with stacked generative adversarial networks, 2018.
- [6] Zizhao Zhang, Yuanpu Xie, and Lin Yang. Photographic text-to-image synthesis with a hierarchically-nested adversarial network, 2018.
- [7] Tingting Qiao, Jing Zhang, Duanqing Xu, and Dacheng Tao. MirrorGAN: Learning text-to-image generation by redescription, 2019. 1
- [8] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Raphael Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Imagen: Photorealistic text-to-image diffusion models, 2022. Imagen. 1, 2
- [9] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation, 2021. 1
- [10] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Dall-e: Creating images from text, 2021. OpenAI. 1, 2
- [11] MidJourney, Inc. MidJourney, 2022. Accessed: 2025-02-27. 1
- [12] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. Stable Diffusion. 1, 2
- [13] DeepFloyd. DeepFloyd if: Text-to-image generation model, 2023. DeepFloyd. 1
- [14] Oran Gafni, Lior Reichart, Yuval Atzmon, Eliya Nachmani, Raja Giryes, Lior Wolf, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. <https://make-a-scene.github.io>, 2022. Make-A-Scene.
- [15] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Junyuan Shang, Yujie Qian, Qiankun Zhao, Weilin Zhao, Zilong Zheng, Zhou Shao, Jiazhan Feng, Jianping He, Xinghan Liu, Yifan Xu, and Jie Tang. Cogview: Mastering text-to-image generation via transformers. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 2285–2294. ACM, 2021. CogView. 1
- [16] Kihyuk Sohn, Nataniel Ruiz, Kimin Lee, Daniel Castro Chin, Irina Blok, Huiwen Chang, Jarred Barber, Lu Jiang, Glenn Entis, Yuanzhen Li, Yuan Hao, Irfan Essa, Michael Rubinstein, and Dilip Krishnan. Styledrop: Text-to-image generation in any style, 2023. 2
- [17] 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, 2023. 2
- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2019. 2
- [19] Tobias Kuhn, Steven Bourke, Levin Brinkmann, Tobias Buchwald, Conor Digan, Hendrik Hache, Sebastian Jaeger, Patrick Lehmann, Oskar Maier, Stefan Matting, and Yura Okulovsky. Supporting stylists by recommending fashion style, 2019. 2
- [20] Amir Hertz, Andrey Voynov, Shlomi Fruchter, and Daniel Cohen-Or. Style aligned image generation via shared attention. In *arXiv preprint arxiv:2312.02133*, 2023. 2, 3, 5, 6, 7
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan, 2020. 2, 4, 6, 7
- [22] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021. 2, 5
- [23] 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 (ICCV)*, pages 1501–1510, 2017. 3
- [24] Ken Shoemake. Animating rotation with quaternion curves. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, pages 245–254, 1985. 4
- [25] Stability AI. Stable diffusion xl - base 1.0, 2023. Accessed: September 2024. 6
- [26] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)*, 2021. 6
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 6
- [28] Dongxu Li, Junnan Li, and Steven C. H. Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing, 2023. 6
- [29] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision, 2022. 6
- [30] Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. Clap: Learning audio concepts from natural language supervision, 2022. 6
- [31] OpenWeather. Openweather api, 2024. Accessed: 2024-10-31. 6

- [32] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. [6](#)