

EmoStyle: Emotion-Driven Image Stylization

Jingyuan Yang, Zihuan Bai, Hui Huang*

CSSE, Shenzhen University

{jingyuanyang.jyy, baizihuan23, hhzhiyan}@gmail.com



Figure 1. Affective Image Stylization with EmoStyle, aiming to transform user-provided images through artistic stylization to evoke specific emotional responses. It requires only emotion words as prompts, eliminating the need for reference images or detailed text descriptions.

Abstract

Art has long been a profound medium for expressing emotions. While existing image stylization methods effectively transform visual appearance, they often overlook the emotional impact carried by styles. To bridge this gap, we introduce Affective Image Stylization (AIS), a task that applies artistic styles to evoke specific emotions while preserving content. We present EmoStyle, a framework designed to address key challenges in AIS, including the lack of training data and the emotion–style mapping. First, we construct EmoStyleSet, a content-emotion-stylized image triplet dataset derived from ArtEmis to support AIS. We then propose an Emotion–Content Reasoner that adaptively integrates emotional cues with content to learn coherent style queries. Given the discrete nature of artistic styles, we further develop a Style Quantizer that converts continuous style features into emotion-related codebook entries. Extensive qualitative and quantitative evaluations, including user studies, demonstrate that EmoStyle enhances emotional expressiveness while maintaining content consistency. Moreover, the learned emotion-aware style dictionary is adaptable to other generative tasks, highlighting its potential for broader applications. Our work establishes

*Corresponding author

a foundation for emotion-driven image stylization, expanding the creative potential of AI-generated art. Project page: <https://vcc.tech/research/2026/EmoStyle>.

1. Introduction

“Art is not a handicraft, it is the transmission of feelings.”

—Leo Tolstoy

Beyond aesthetic appreciation, art has long served as a powerful medium for conveying emotions, allowing artists to elicit deep emotional responses from viewers [2, 18]. For instance, while *Van Gogh’s Starry Night* evokes awe and wonder, *Edvard Munch’s The Scream* conveys fear and despair. These expressive contrasts raise a fundamental question: *Can we automatically create images that are not only visually compelling but also emotionally resonant?*

Style transfer (ST) applies artistic techniques to modify an image’s appearance while preserving its structure, which has advanced significantly in recent years [8, 14, 37, 41]. However, most methods rely on reference images, limiting their applicability when no style image is available. Text-guided style transfer [15, 20, 26] offers greater flexibility by allowing users to specify styles through prompts. Yet, current methods usually describe style via medium (*oil painting*), artist (*Claude Monet*), art movement (*Cubism*),

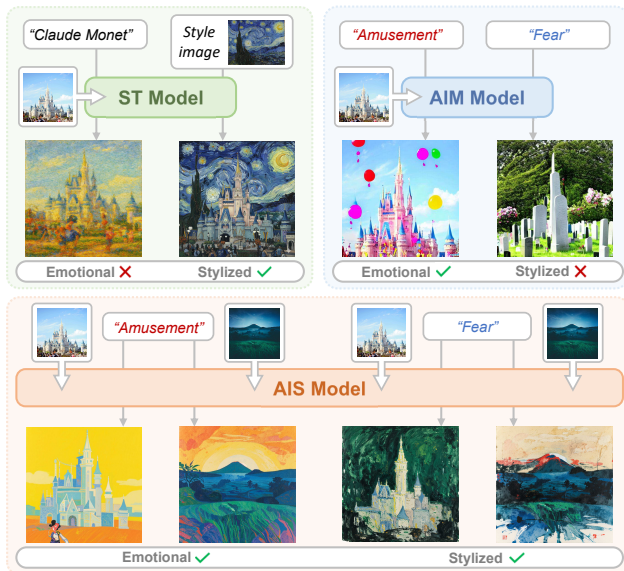


Figure 2. Comparison among Style Transfer (ST), Affective Image Manipulation (AIM), and our proposed Affective Image Stylization (AIS). While ST focuses on style and AIM emphasizes emotion, AIS generates emotionally expressive and stylized results.

or iconic artworks (*Mona Lisa*), necessitating a certain degree of artistic knowledge. While not everyone is an artist, emotions are universally understood.

In Affective Computing [31], researchers have extensively studied how images convey emotions [1, 3, 7, 52, 53]. With the rise of generative AI, the focus has shifted from understanding emotions to Affective Image Manipulation (AIM), which modifies images to evoke target emotional responses [44]. Some methods alter emotional tones through color adjustments [25, 36, 42], while others modify visual content [23, 43, 44]. However, these techniques primarily focus on generating realistic images, overlooking artistic styles as expressive tools for emotional communication. Recent studies have explored emotion-aware style transfer [34, 39], yet they often rely on reference images [34] or require carefully crafted textual descriptions [39], limiting their practicality and accessibility.

Therefore, we introduce Affective Image Stylization (AIS), a new task that integrates artistic style into images to evoke specific emotional responses while preserving semantic content. As shown in Fig. 2, AIS bridges the gap between ST and AIM by combining the strengths of both: it evokes emotions like AIM while producing stylized results like ST, prompted only with emotion words. AIS introduces two key challenges: (1) *Lack of training data*: no dataset provides content-emotion-stylized image triplets for learning emotion-aware stylization; (2) *Emotion-style mapping*: generating stylistic variations that remain content-consistent yet emotionally expressive is highly non-trivial.

We therefore introduce both a tailored dataset and a framework to address these challenges.

AIS requires the model to understand *what* the image depicts (content) and *which* emotion it should express, in order to determine *how* it should appear (style). Thus, AIS forms a content-emotion-stylized image triplet problem, and we construct EmoStyleSet from ArtEmis [1] to support it. Stylized images and emotion labels come directly from ArtEmis. To obtain content images, we adopt UnZipLoRA [24] to disentangle content from style and use ControlNet [47] to generate realistic images. Automatic and human evaluations are applied to filter noisy results.

We propose EmoStyle, an AIS framework capable of generating emotion-evoking, content-preserving images with appealing artistic styles, see Fig. 1. EmoStyle consists of two main modules: *Emotion-Content Reasoner* and *Style Quantizer*. In art, style and content are intertwined, with stylistic choices influenced by subject matter and emotion [12]. Inspired by this, we develop an Emotion-Content Reasoner that integrates emotion and content cues through a transformer [19] to learn coherent style queries. Unlike content, style often falls into discrete perceptual categories, e.g., *Impressionism*, *Modernism*, and *Realism*. We therefore introduce a Style Quantizer that discretizes continuous features into codebook entries, each representing a style prototype associated with a specific emotion. This enables the Style Quantizer to align each emotion with multiple artistic styles, allowing for controllable and interpretable AIS.

To evaluate EmoStyle, we use the 405-image inference set following EmoEdit, given the same input format, i.e., emotion label and content image. Beyond test images, our method applies to any user-provided content. We compare EmoStyle with state-of-the-art style transfer, image editing, and AIM approaches. Both quantitative and qualitative results confirm its effectiveness, further supported by a user study validating its alignment with human preferences. Additionally, the learned style dictionaries show strong adaptability to emotion-driven text-to-image generation tasks.

In summary, our contributions are:

- We introduce Affective Image Stylization (AIS), which applies artistic styles to evoke target emotions while preserving its original content, and construct EmoStyleSet as its data foundation.
- We propose EmoStyle, which integrates an Emotion-Content Reasoner to fuse emotion and content and a Style Quantizer to map continuous features into discrete, emotion-aware style representations.
- We conduct extensive experiments and user studies, showing that EmoStyle achieves a strong balance between emotional expressiveness and content consistency. It also shows strong potential for affective art creation.

2. Related work

2.1. Artwork Emotion Analysis

As a powerful medium for expressing emotions, art has garnered increasing research interest in Artwork Emotion Analysis in recent years [1, 3, 7, 16, 28, 46]. Leveraging the large-scale WikiArt dataset [33], ArtEmis [1] and ArtEmis-V2 [28] enriched emotion-related attributes and employed a trained neural speaker to capture finer emotional nuances in artworks. Joshi *et al.* [18] analyzed the problems of aesthetics and emotions in images from a computational perspective and discussed future directions in this research domain. Chen *et al.* [7] investigated local emotion stimuli and established a benchmark dataset with baseline evaluations. Zhang *et al.* [46] introduced EmoArt, a large-scale art dataset where each image is carefully annotated with fine-grained emotion labels. Nevertheless, previous studies have primarily focused on recognizing and interpreting emotions in artworks rather than generating them. With the rise of generative AI, we introduce Affective Image Stylization (AIS) to explore how artistic styles can be used to evoke visual emotions.

2.2. Style Transfer

The task of style transfer can be categorized into reference image-based style transfer [6, 8, 11, 14, 37, 38, 40, 41] and text-guided style transfer [10, 15, 20, 21, 37]. Among image-based methods, Wang *et al.* [38] introduced a new method to achieve disentangle content-style in the CLIP image space. Wu *et al.* [40] presented a well-performed model via disentangled and reward learning. In text-guided approaches, Kwon *et al.* [20] proposed a framework that enables style transfer solely through text descriptions. Fu *et al.* [10] introduced Contrastive Language-Visual Artist (CLVA), which learns to extract visual semantics from style instructions using a patch-wise style discriminator. Zhang *et al.* [50] treated style as a learnable textual description, effectively capturing and transferring artistic styles. Wang *et al.* [37] accomplished both text-guided and image-based tasks by constructing a high-quality dataset. While existing methods can produce aesthetically appealing results, they often rely on style images or art-knowledge prompts, which demands a certain level of expertise. Emotions, in contrast, are universally understood and easy to specify. To bridge this gap, we use emotion words as prompts to guide stylized generation and construct a dataset to support this task.

2.3. Affective Image Manipulation

Affective Image Manipulation (AIM) has evolved along two directions: style-based methods, which modify colors and styles to evoke emotions [25, 34, 39, 42, 48], and content-based methods, which incorporate emotional elements into images [23, 44, 49]. Early AIM research primarily fo-

cused on color manipulation to influence emotions. Yang *et al.* [42] pioneered the use of color transfer for emotion-driven image editing, while Liu *et al.* [25] introduced an end-to-end color transfer approach. More recently, AIM has expanded beyond color adjustments to incorporate emotional elements into image content. Yang *et al.* [44] introduced EmoEdit, building a task-specific dataset and training an Emotion Adapter to modify image content to express emotions. Zhang *et al.* [49] designed an emotional mapper and utilized an MLLM for supervised training. Lin *et al.* [23] proposed EmoEditor, a human-preference aligned diffusion model, which evokes emotions by manipulating image regions. Meanwhile, style-based AIM has evolved to explore emotion-aware stylization. Sun *et al.* [34] developed MSNet, a deep architecture that employs a Multi-Sentiment Semantics Space for sentiment-aware style transfer. Weng *et al.* [39] introduced the Affective Image Filter (AIF), designing to translate textual descriptions into visual representations. Zhang *et al.* [48] leveraged the generative priors of diffusion models and proposed an extension of AIF. Despite significant progress in emotion-aware style transfer, existing methods often rely on reference images or carefully crafted textual descriptions, which limits their accessibility. Given an emotion word and a content image, EmoStyle introduces an Emotion-Content Reasoner to extract emotion-aware style query and a Style Quantizer to fetch its suitable prototype. The style dictionary establishes an interpretable mapping between emotions and styles, ensuring the expressive generative results.

3. Method

3.1. EmoStyleSet

Existing artwork datasets (*e.g.*, ArtEmis [1], EmoArt [46]) serve as valuable benchmarks for artwork emotion analysis but remain inadequate for Affective Image Stylization (AIS). These datasets assign emotion labels to artworks without distinguishing whether the emotion stems from *content* or *style*. Since AIS aims to learn how style conveys emotion, separating style from overall artwork is essential. To this end, we build EmoStyleSet, derived from ArtEmis, containing 10,041 triplets of content image, target emotion, and stylized image.

Specifically, ArtEmis is a large-scale visual emotion dataset containing 80K artworks, designed to capture the affective responses elicited by visual art. Each image is annotated by multiple evaluators, resulting in an emotion distribution per artwork. We convert these distributions into single labels and remove neutral samples to obtain simpler and accurate emotion annotations. As both content and style influence emotional perception [17, 29], we separate content from style to study how artistic styles evoke emotions.

To extract the content of each artwork, we employ BLIP-

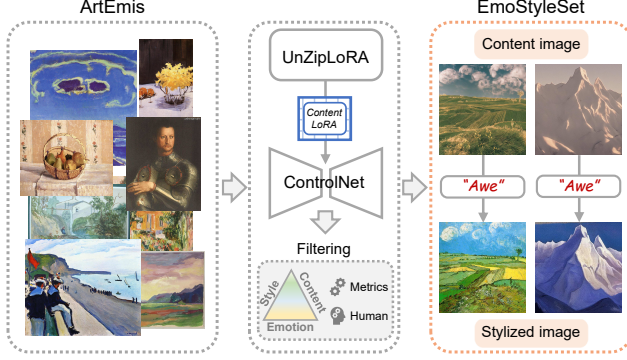


Figure 3. Construction process of EmoStyleSet. Given artworks from ArtEmis, after generation and filtering, each triplet contains a content image, a target emotion, and a stylized image.

2 [22] to generate descriptive captions for ArtEmis. Given an artwork and its caption, UnZipLoRA [24] decomposes the image into a content LoRA and a style LoRA. We further convert each artwork into a Canny edge map to preserve structural information. The Canny image and caption are then jointly input to ControlNet [47], where the content LoRA is injected to maintain semantic consistency. The output of ControlNet is regarded as the content image, with the style removed from the original ArtEmis artwork.

Since the content images in EmoStyleSet are generated in an unsupervised manner, certain inaccuracies may occur. To ensure data quality, we perform filtering from three aspects: style, content and emotion, as illustrated in Fig. 3. For content, we evaluate semantic consistency using CLIP similarity [32] and measure structural preservation with LPIPS. Emotion accuracy is assessed by a classifier trained on ArtEmis, ensuring each triplet matches its target emotion. Finally, human annotators verify that the stylized images display clear stylistic differences from their corresponding content images. More details are provided in the supplementary materials.

As shown in Fig. 3, EmoStyleSet is organized as {content image, emotion word, stylized image}, denoted as content-emotion-stylized image triplets, following the standard style transfer paradigm. Taking *awe* as an example, the upper row shows content images of landscapes and mountains, while the lower row presents their stylized counterparts that evoke a sense of magnificence and vastness.

3.2. EmoStyle

3.2.1. Emotion-Content Reasoner

In artistic creation, style and content are deeply intertwined, with artists naturally adapting their stylistic choices to align with the subject matter and emotional intent [12]. Consequently, a key challenge in AIS is emotion-aware stylization, *i.e.*, evoking emotions through appropriate artistic

styles. To address it, we first propose Emotion-Content Reasoner, designed to determine the most appropriate style given a content image and a target emotion. As illustrated in Fig. 4, we employ SigLIP [45] to extract semantic features from the content image. Prior works [43, 44] modeled emotions as textual inputs, which is inadequate because LLMs often associate emotion words with facial expressions, whereas emotions in AIS involve a broader artistic context. Therefore, we encode each emotion as a 1*8 one-hot vector, which offers two advantages: (1) emotions are mutually orthogonal; (2) all vectors collectively span the emotion space. Since semantic and emotion features lie in different spaces, we project them through separate layers into a unified embedding space. The fused embeddings are then processed by a four-block multi-modal transformer [19], which applies both self- and cross-attention to model interactions between emotion and content features. Through this interaction, the transformer reasons across modalities to derive an emotion-aware, content-conditioned style query q_i for subsequent style selection:

$$\begin{aligned} \overline{q}_i^k &= MSA(LN(q_i^{k-1})) + q_i^{k-1}, \\ q_i^k &= MLP(LN(\overline{q}_i^k)) + \overline{q}_i^k, \end{aligned} \quad (1)$$

where LN denotes the LayerNorm, k indexes the transformer blocks. Specifically, we initialize q_i^0 by concatenating the outputs of emotion and content projections.

3.2.2. Style Quantizer

Artistic styles are often perceived categorically rather than continuously in human visual cognition, *e.g.*, Impressionism, Modernism, Realism. Motivated by this, we introduce a Style Quantizer to discretize style features to distinct prototypes, enabling interpretable and controllable style learning. The relationship between emotion and style is inherently *many-to-many*, as each emotion can be expressed through various styles, and each style can evoke multiple emotions. However, for the initial exploration in AIS, we adopt a *one-to-many* simplification, following art history that emotions are conveyed through diverse styles [13].

In EmoStyleSet, each triplet is denoted as $\{I_c, e, I_s\}$, where I_c represents content image, e refers to emotion word, and I_s is stylized image. Inspired by VQ-VAE [35], we construct eight style dictionaries, each corresponding to an emotion category, denoted as $Z^e = \{z_k^e\}_{k=1}^K$. To be specific, we employ the style encoder $\mathcal{E}_s(\cdot)$ from USO [40] to extract style features and compute pairwise similarities among all artworks. A similarity threshold is then applied to initialize each style dictionary with distinct style prototypes. In the first stage, we learn the prototype styles using Eq. 3. In the second stage, we apply vector quantization [35] $Q(\cdot)$ to the style query q_i to fetch its nearest neighbor z_k^e from style dictionary Z^e , thereby converting the representation from continuous to discrete:

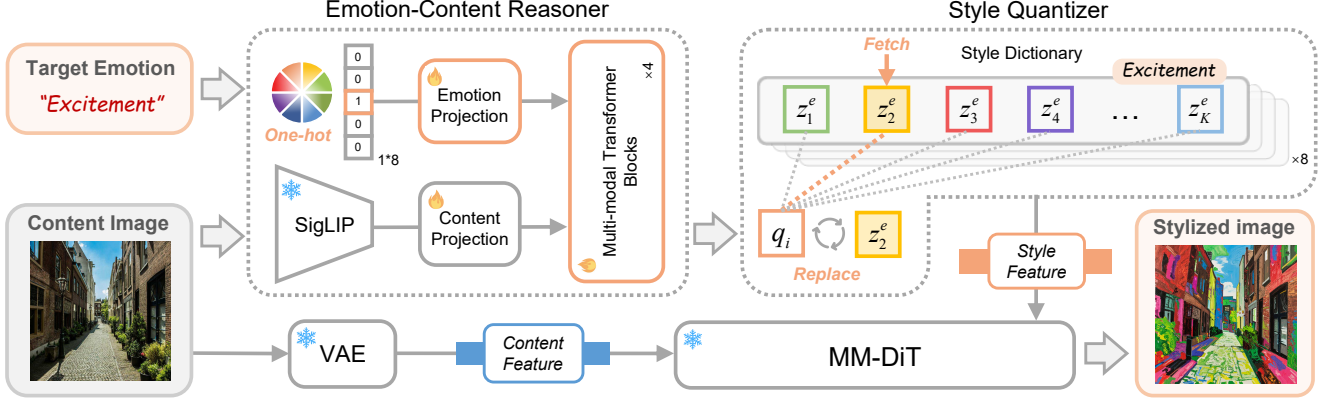


Figure 4. Overview of EmoStyle. We introduce an Emotion–Content Reasoner to integrate emotion and content features, and a Style Quantizer to map continuous queries to discrete style prototypes, generating stylized images with faithful emotion and preserved content.

$$Q(q_i) = z_k^e, \quad \text{where } k = \arg \min_j \|q_i - z_j^e\|_2. \quad (2)$$

By adopting discrete style representations, we simplify style-emotion mapping and ensure coherence between style and content in artworks, effectively addressing the challenge outlined in the introduction.

3.2.3. Training and Inference

As shown in Fig. 4, we adopt MM-DiT [40] as the backbone. Specifically, the content image is encoded by a VAE into a latent content feature, while the output of the Style Quantizer serves as the style feature. These two features are then fed into MM-DiT for generation. During training, we freeze the MM-DiT parameters and update only the Emotion–Content Reasoner and Style Quantizer. Style Quantizer learns what each style prototype represents, while Emotion–Content Reasoner learns which prototype to select given an emotion–content pair. Accordingly, EmoStyle adopts a two-stage training scheme: prototype learning followed by emotion–style alignment.

In the first stage, we learn the prototype style features Z^e with the stylized images from EmoStyleSet:

$$\begin{aligned} \mathcal{L}_{style} &= \|z_k^e - \mathcal{E}_s(I_s)\|_2^2, \\ \text{where } k &= \arg \min_j \|\mathcal{E}_s(I_s) - z_j^e\|_2, \end{aligned} \quad (3)$$

where z_k^e is the k -th representation in the style dictionary for emotion e . This loss guides the style representations to form discrete and diverse prototypes distributed across the style space. It can be viewed as a clustering process, where each style feature is assigned to its nearest cluster centroid, and the centroids are iteratively updated based on their assigned members. As a result, each style dictionary

contains representative style embeddings for each emotion, as visualized in Sec. 4.

In the second stage, we use content-emotion-stylized image triplets to align the generated stylized results with the ground truth at both the feature and pixel levels. To ensure high-quality generation and effectively guide pixel-level learning on paired data, we adopt the standard flow-matching loss:

$$\mathcal{L}_{FM} = \mathbb{E}_{x_0, t, \epsilon} [w(t) \|v_\theta - v_t\|^2], \quad (4)$$

where v_θ denotes the predicted velocity and $v_t = \frac{d\alpha_t}{dt} x_0 + \frac{d\sigma_t}{dt} \epsilon$ the ground-truth velocity.

Moreover, for a given emotion and content image pair, to guide emotion-content reasoner to find its corresponding style representation, we propose an align loss:

$$\mathcal{L}_{align} = \|q_i - z_k^e\|_2^2, \quad (5)$$

where z_k^e is the nearest neighbor of q_i , ensuring that the learned style queries follow the style distribution in EmoStyleSet.

Besides, to ensure emotion fidelity, we use the emotion score e_n derived from the voting results in ArtEmis to weight the training losses of different samples:

$$\begin{aligned} \mathcal{L}_1 &= \frac{1}{N} \sum_n e_n \cdot \mathcal{L}_{style}, \\ \mathcal{L}_2 &= \frac{1}{N} \sum_n e_n \cdot (\mathcal{L}_{FM} + \mathcal{L}_{align}), \end{aligned} \quad (6)$$

where e_n denotes the emotion score of the n -th sample and N refers to the total number of training samples.

At inference, once the user provides a content image (e.g., a street view) and a target emotion (e.g., excitement), EmoStyle can generate a stylized image that is both emotion-evoking and aesthetically pleasing, as in Fig. 4.

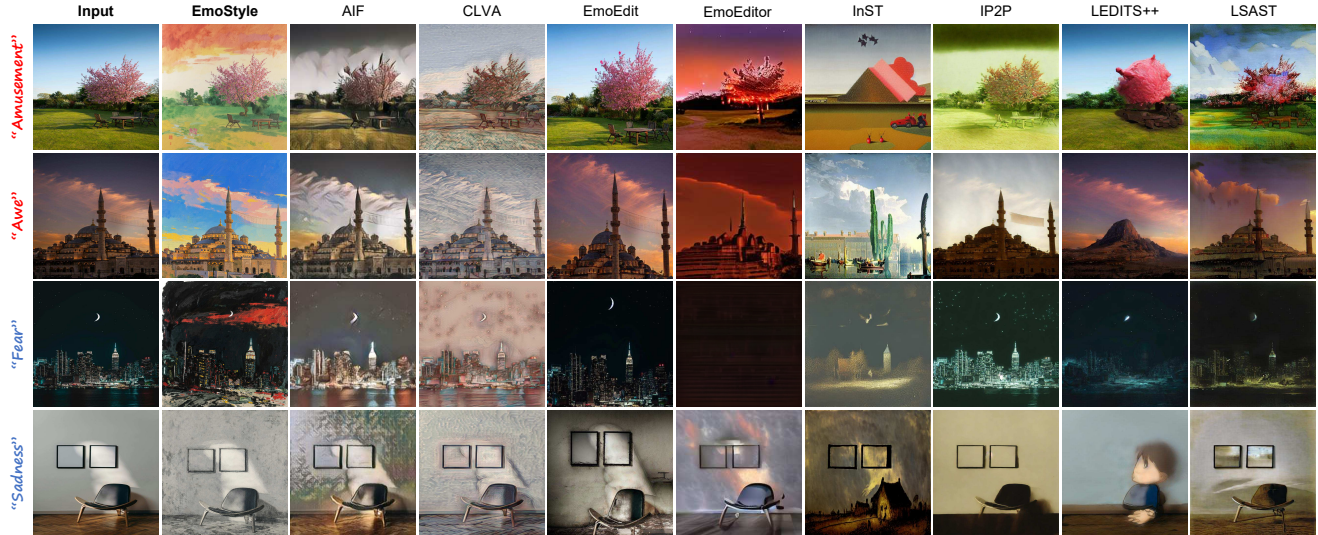


Figure 5. Comparison with the state-of-the-art methods, where EmoStyle surpasses others on emotion fidelity and aesthetic appeal.

4. Experiments

4.1. Dataset and Evaluation

Dataset We train EmoStyle using content-emotion-stylized image triplets from the curated EmoStyleSet, as detailed in Sec. 3.2.3. Since EmoStyle shares the same input format with EmoEdit [44], a content image and an emotion category, we evaluate it on the EmoEdit inference set, which contains 405 realistic, user-uploaded images collected from the internet. Unlike EmoEdit, which produces realistic outputs, EmoStyle performs AIS and generates stylized results. Each image is stylized into eight target emotions from Mikels’ wheel [27]: *amusement*, *awe*, *contentment*, *excitement*, *anger*, *disgust*, *fear*, and *sadness*. In total, our inference set contains 3,240 stylized results.

Evaluation Metrics Since AIS is a new task, we compile a set of metrics to evaluate it from multiple perspectives. Like most style transfer methods, we first evaluate content preservation across different approaches. Following OmniStyle [37], we employ CLIP [32] to assess semantic consistency between content and stylized images, and DINO [30] to measure pairwise structural preservation. Unlike conventional style transfer methods, AIS focuses on emotional expressiveness in stylized images. Thus, we introduce two emotion-related metrics. Sentiment Gap (SG) [39] measures how effectively the stylized images evoke the intended emotional responses. Emotion Accuracy (Emo-A) assesses whether the stylized image aligns with the target emotion, using a pre-trained emotion classifier from ArtEmis [1]. In addition, Style Difference (SD) [9] evaluates whether the stylized images exhibit appropriate

Table 1. Comparisons with the state-of-the-art methods on style transfer, image editing and AIM methods.

Method	CLIP \uparrow	DINO \uparrow	SG (%) \downarrow	Emo-A (%) \uparrow	SD \downarrow
LSAST [51]	0.551	0.747	2.231	12.50	11.28
CLIPStyler [20]	0.709	0.769	3.001	12.60	19.89
InST [50]	0.569	0.679	<u>2.016</u>	21.22	11.48
OmniStyle [37]	0.710	<u>0.813</u>	2.615	12.80	11.90
IP2P [5]	0.708	0.729	3.459	<u>24.34</u>	12.76
LEDITS++ [4]	0.687	0.807	2.637	15.97	13.11
EmoEditor [23]	0.686	0.761	2.744	14.88	13.65
EmoEdit [44]	0.597	0.545	2.245	12.60	28.83
CLVA [10]	0.727	0.789	2.030	14.99	9.49
AIF [39]	0.712	0.780	2.625	12.99	<u>8.48</u>
EmoStyle	<u>0.718</u>	0.842	1.976	33.36	7.59

color and textual consistency with the training data. For more details, please refer to the supplementary materials.

4.2. Comparisons

To validate the effectiveness of EmoStyle, we compare it against the state-of-the-art methods across three categories: (1) Style transfer: LSAST [51], CLIPStyler [20], InST [50], and OmniStyle [37]; (2) Image editing: LEDITS++ [4], and IP2P [5]; (3) AIM: EmoEditor [23], EmoEdit [44], CLVA [10], and AIF [39].

Quantitative Comparisons As shown in Table 1, EmoStyle outperforms other methods across most metrics, effectively evoking emotions while preserving content. Unlike existing approaches that fail to capture emotional knowledge, EmoStyle achieves the highest Emo-A (33.36%), surpassing the second-best by nearly 9%, and records the lowest SG (1.976), owing to the effective inter-




Figure 6. Ablation study. The Emotion Encoder, Emotion-Content Reasoner, and Style Quantizer are all shown to be effective.

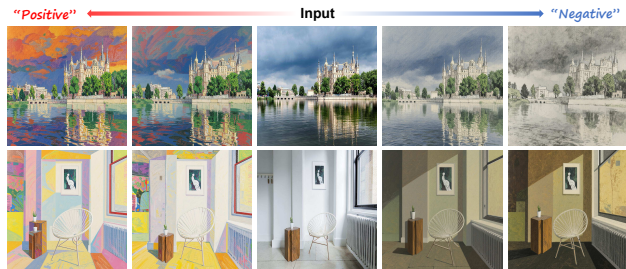


Figure 7. Ablation study on image guidance scale. EmoStyle can progressively edit an image towards different emotional polarities.

action in the Emotion-Content Reasoner. It also ranks first in DINO (0.842) and second in CLIP (0.718), demonstrating strong structural and semantic preservation with expressive styles. While InST evokes emotions better than most methods, *i.e.*, lower SG, it struggles to preserve semantic information. The SD results further validate EmoStyle’s Style Quantizer, confirming its ability to learn accurate style representations from EmoStyleSet. Overall, these results demonstrate EmoStyle’s superior balance between emotion expression and content preservation, making it an effective framework for emotion-driven artistic stylization.

Qualitative Comparisons We present the qualitative results in Fig. 5. Most compared methods struggle to convey emotions through stylization due to the lack of emotional knowledge. Style transfer methods such as LSAST and InST can transform content images into artworks but fail to evoke the target emotion, resulting in visually appealing yet emotionally ambiguous results. Among the two image editing methods, IP2P effectively introduces artistic elements and can evoke emotions to some extent after being trained on EmoStyleSet, highlighting the quality of the proposed dataset. AIM methods such as EmoEdit and EmoEditor primarily modify colors or add emotional semantics to evoke emotions, but fail to achieve effective stylization. Both CLVA and AIF can generate stylized images, but the style differences across emotions remain subtle. By explicitly modeling the relationship between emotions and artistic




Figure 8. Visualization of the style dictionary, where each column corresponds to a distinct prototype z_k^e for a specific emotion.

Table 2. User preference study. The numbers indicate the percentage of participants who vote for the result.

Method	Aesthetic Perception \uparrow	Emotion fidelity \uparrow	Balance \uparrow
CLVA [10]	8.50 \pm 12.13%	0.81 \pm 2.16%	1.19 \pm 7.07%
InST [50]	2.50 \pm 4.86%	29.63 \pm 2.56%	1.34 \pm 5.53%
AIF [39]	9.08 \pm 9.42%	5.09 \pm 2.22%	7.76 \pm 9.91%
EmoStyle	79.92\pm21.24%	64.47\pm4.55%	89.70\pm14.48%

styles, our approach produces stylized results that are both aesthetically appealing and emotionally evocative.

User Study Since emotional perception is inherently abstract and subjective, we conducted a user study to assess human preferences for our method. The study presented 24 sets of images, each containing one content image and four stylized versions generated by CLVA, InST, AIF and EmoStyle. Participants were shown image sets and asked three questions regarding emotion, aesthetics, and their balance. As shown in Table 2, EmoStyle is the most preferred method across all evaluation aspects. Balancing emotion and aesthetics is difficult, yet EmoStyle achieved the highest scores in both aspects, *i.e.*, 64.47% and 79.92%, respectively. Moreover, when evaluating the balance between emotion and aesthetics, 89.70% of participants favored EmoStyle, demonstrating its strong alignment with human perception. InST ranked the second in evoking emotions, scoring 29.63%. However, although it introduces emotional styles to images, it fails to preserve the original structure and semantics, as illustrated in Fig. 5.

4.3. Ablation Study

Methodology In Fig. 6, we evaluate the effectiveness of the key components in EmoStyle, namely the Emotion Encoder, Emotion-Content Reasoner, and Style Quantizer. Without the Style Quantizer, the results tend to appear overly realistic, indicating that the Emotion-Content Reasoner alone cannot map emotions to expressive artistic styles. The Emotion Encoder and Emotion-Content Reasoner are both essential for emotion-aware stylization, as removing them leads to results that are less emotion-evoking. Overall, through the integration of these key components,



Figure 9. Emotion-driven text-to-image generation. Conditioned on a text prompt (content) and an emotion word (style), EmoStyle produces semantically faithful, emotionally expressive and aesthetically appealing stylized results.

EmoStyle generates emotionally faithful, artistically appealing, and content-consistent images.

Guidance Scale We present EmoStyle’s stylized results under varying image guidance scale in Fig. 6. The middle image is the input, while the sides show stylization under varying guidance scales for positive and negative emotions. Increasing the scale enhances emotion and style expressiveness but compromises structural preservation, revealing a trade-off between emotion and content. However, EmoStyle effectively balances these aspects, as demonstrated in Table 1, Table 2, and Fig. 5. In addition, EmoStyle offers users flexible control over the image guidance scale, enabling adjustment of emotion intensity and fine-tuning of stylization according to personal preference.

Visualization As shown in Fig. 6, we visualize several style representations in the Style dictionary $Z^e = \{z_k^e\}_{k=1}^K$ for two distinct emotions. Each style dictionary offers diverse and aesthetically appealing prototypes that effectively evoke the target emotion, reflecting the one-to-many emotion style mapping assumed in Sec. 3.2.2. Users can further choose specific prototypes for each emotion to generate images that match their preferences.

4.4. Applications

Emotion-driven Text-to-Image Generation Beyond image stylization, EmoStyle can be extended to text-to-image generation, enabling the creation of emotionally expressive images from textual descriptions, as shown in Fig. 9. Since SigLIP aligns image and text modalities, we replace the image encoder in Fig. 4 with a text encoder for this task. The Emotion-Content Reasoner then selects the most appropriate emotion-aware style based on the given textual description. For example, given the prompt “Dog”, our method generates eight stylized images, each conveying a target emotion while remaining artistically appealing. These results highlight EmoStyle’s ability to preserve content in-

tegrity while ensuring emotion fidelity, expanding its potential for affective content creation and personalized design.

5. Conclusion

Discussion We present EmoStyle, a framework that applies artistic styles to evoke specific emotions while preserving content integrity. Unlike conventional stylization methods that prioritize aesthetics, EmoStyle focuses on emotion-driven transformations. EmoStyle ensures content-consistent, emotion-aware stylization, and offers multiple style prototypes for each emotion. Experiments have demonstrated that EmoStyle effectively generates emotionally expressive and aesthetically appealing images. As the initial attempt to explore the intrinsic relationship between emotion and style, EmoStyle offers greater flexibility and potential for broader applications in visual art design.

Limitations Despite its advancements, EmoStyle still faces several challenges. First, the mapping between emotion and style is inherently many-to-many. Dark, muted tones may evoke melancholy in one image but mystery in another, while a single emotion can be expressed through diverse artistic styles ranging from watercolor to impressionism. Second, emotion in artworks is evoked not only by style but also by content. Understanding how the two visual stimuli interact to elicit emotions and developing methods that balance their contributions remain open research questions. Finally, although our evaluations capture several aspects, assessing the subjective emotional impact of stylized images is still challenging. Future work should establish more comprehensive evaluation frameworks that incorporate human feedback and psychological insights.

Acknowledgments: This work was supported in part by ICFCRT (W2441020), NSFC (62302312), Guangdong Science and Technology Program (2023A1515011440), Shenzhen Science and Technology Program (KQTD20210811090044003, RCJC20200714114435012), and Scientific Development Funds from Shenzhen University.

References

- [1] Panos Achlioptas, Maks Ovsjanikov, Kilichbek Haydarov, Mohamed Elhoseiny, and Leonidas J Guibas. Artemis: Affective language for visual art. In *CVPR*, pages 11569–11579, 2021. 2, 3, 6
- [2] Ismay Barwell. How does art express emotion? *The Journal of Aesthetics and Art Criticism*, 45(2):175–181, 1986. 1
- [3] Digbalay Bose, Krishna Somandepalli, Souvik Kundu, Rimita Lahiri, Jonathan Gratch, and Shrikanth Narayanan. Understanding of emotion perception from art. *arXiv preprint arXiv:2110.06486*, 2021. 2, 3
- [4] Manuel Brack, Felix Friedrich, Katharia Kornmeier, Linoy Tsaban, Patrick Schramowski, Kristian Kersting, and Apolinário Passos. Ledits++: Limitless image editing using text-to-image models. In *CVPR*, pages 8861–8870, 2024. 6
- [5] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, pages 18392–18402, 2023. 6
- [6] Haibo Chen, Lei Zhao, Zhizhong Wang, Huiming Zhang, Zhiwen Zuo, Ailin Li, Wei Xing, and Dongming Lu. Dualast: Dual style-learning networks for artistic style transfer. In *CVPR*, pages 872–881, 2021. 3
- [7] Tianwei Chen, Noa Garcia, Liangzhi Li, and Yuta Nakashima. Exploring emotional stimuli detection in artworks: A benchmark dataset and baselines evaluation. *Journal of Imaging*, 10(6):136, 2024. 2, 3
- [8] 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 *CVPR*, pages 8795–8805, 2024. 1, 3
- [9] Yingying Deng, Fan Tang, Weiming Dong, Chongyang Ma, Xingjia Pan, Lei Wang, and Changsheng Xu. Stytr2: Image style transfer with transformers. In *CVPR*, pages 11326–11336, 2022. 6
- [10] Tsu-Jui Fu, Xin Eric Wang, and William Yang Wang. Language-driven artistic style transfer. In *ECCV*, pages 717–734, 2022. 3, 6, 7
- [11] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, pages 2414–2423, 2016. 3
- [12] Ernst Hans Gombrich and EH Gombrich. *The story of art*. 1995. 2, 4
- [13] Charles Harrison and Paul Wood. *Art in theory, 1900–2000 an anthology of changing ideas*. 2003. 4
- [14] Amir Hertz, Andrey Voynov, Shlomi Fruchter, and Daniel Cohen-Or. Style aligned image generation via shared attention. In *CVPR*, pages 4775–4785, 2024. 1, 3
- [15] Nisha Huang, Yuxin Zhang, Fan Tang, Chongyang Ma, Haibin Huang, Weiming Dong, and Changsheng Xu. Diffstyler: Controllable dual diffusion for text-driven image stylization. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 1, 3
- [16] Shintaro Ishikawa and Komei Sugiura. Affective image captioning for visual artworks using emotion-based cross-attention mechanisms. *IEEE Access*, 11:24527–24534, 2023. 3
- [17] Jia Jia, Sen Wu, Xiaohui Wang, Peiyun Hu, Lianhong Cai, and Jie Tang. Can we understand van gogh’s mood? learning to infer affects from images in social networks. In *ACM MM*, pages 857–860, 2012. 3
- [18] Dhiraj Joshi, Ritendra Datta, Elena Fedorovskaya, Quang-Tuan Luong, James Z Wang, Jia Li, and Jiebo Luo. Aesthetics and emotions in images. *IEEE Signal Processing Magazine*, 28(5):94–115, 2011. 1, 3
- [19] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, pages 5583–5594, 2021. 2, 4
- [20] Gihyun Kwon and Jong Chul Ye. Clipstyler: Image style transfer with a single text condition. In *CVPR*, pages 18062–18071, 2022. 1, 3, 6
- [21] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip HS Torr. Manigan: Text-guided image manipulation. In *CVPR*, pages 7880–7889, 2020. 3
- [22] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, pages 19730–19742, 2023. 4
- [23] Qing Lin, Jingfeng Zhang, Yew-Soon Ong, and Mengmi Zhang. Make me happier: Evoking emotions through image diffusion models. In *ICCV*, pages 16367–16376, 2025. 2, 3, 6
- [24] Chang Liu, Viraj Shah, Aiyu Cui, and Svetlana Lazebnik. Unziplora: Separating content and style from a single image. In *CVPR*, pages 16776–16785, 2025. 2, 4
- [25] Da Liu, Yaxi Jiang, Min Pei, and Shiguang Liu. Emotional image color transfer via deep learning. *Pattern Recognition Letters*, 110:16–22, 2018. 2, 3
- [26] Zhi-Song Liu, Li-Wen Wang, Wan-Chi Siu, and Vicky Kalogeiton. Name your style: Text-guided artistic style transfer. In *CVPR Workshops*, pages 3530–3534, 2023. 1
- [27] Joseph A Mikels, Barbara L Fredrickson, Gregory R Larkin, Casey M Lindberg, Sam J Maglio, and Patricia A Reuter-Lorenz. Emotional category data on images from the international affective picture system. *Behavior research methods*, 37:626–630, 2005. 6
- [28] Youssef Mohamed, Faizan Farooq Khan, Kilichbek Haydarov, and Mohamed Elhoseiny. It is okay to not be okay: Overcoming emotional bias in affective image captioning by contrastive data collection. In *CVPR*, pages 21263–21272, 2022. 3
- [29] Saif Mohammad and Svetlana Kiritchenko. Wikiart emotions: An annotated dataset of emotions evoked by art. In *LREC*, 2018. 3
- [30] Maxime Oquab, Timothée Darcet, Theo Moutakanni, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 6
- [31] Rosalind W Picard. *Affective computing*. MIT press, 2000. 2
- [32] 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 *ICML*, pages 8748–8763, 2021. 4, 6

- [33] Babak Saleh and Ahmed Elgammal. Large-scale classification of fine-art paintings: Learning the right metric on the right feature. *arXiv preprint arXiv:1505.00855*, 2015. 3
- [34] Shikun Sun, Jia Jia, Haozhe Wu, Zijie Ye, and Junliang Xing. Msnet: A deep architecture using multi-sentiment semantics for sentiment-aware image style transfer. In *ICASSP*, pages 1–5, 2023. 2, 3
- [35] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *NeurIPS*, 2017. 4
- [36] Xiaohui Wang, Jia Jia, and Lianhong Cai. Affective image adjustment with a single word. *The Visual Computer*, 29: 1121–1133, 2013. 2
- [37] Ye Wang, Ruiqi Liu, Jiang Lin, Fei Liu, Zili Yi, Yilin Wang, and Rui Ma. Omnistyle: Filtering high quality style transfer data at scale. In *CVPR*, pages 7847–7856, 2025. 1, 3, 6
- [38] Zhizhong Wang, Lei Zhao, and Wei Xing. Stylediffusion: Controllable disentangled style transfer via diffusion models. In *ICCV*, pages 7677–7689, 2023. 3
- [39] Shuchen Weng, Peixuan Zhang, Zheng Chang, Xinlong Wang, Si Li, and Boxin Shi. Affective image filter: Reflecting emotions from text to images. In *ICCV*, pages 10810–10819, 2023. 2, 3, 6, 7
- [40] Shaojin Wu, Mengqi Huang, Yufeng Cheng, Wenxu Wu, Ji-ah Tian, Yiming Luo, Fei Ding, and Qian He. Uso: Unified style and subject-driven generation via disentangled and reward learning. *arXiv preprint arXiv:2508.18966*, 2025. 3, 4, 5
- [41] 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. 1, 3
- [42] Chuan-Kai Yang and Li-Kai Peng. Automatic mood-transferring between color images. *IEEE Computer Graphics and Applications*, 28(2):52–61, 2008. 2, 3
- [43] Jingyuan Yang, Jiawei Feng, and Hui Huang. Emogen: Emotional image content generation with text-to-image diffusion models. In *CVPR*, pages 6358–6368, 2024. 2, 4
- [44] Jingyuan Yang, Jiawei Feng, Weibin Luo, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Emoedit: Evoking emotions through image manipulation. In *CVPR*, pages 24690–24699, 2025. 2, 3, 4, 6
- [45] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, pages 11975–11986, 2023. 4
- [46] Cheng Zhang, Hongxia Xie, Bin Wen, Songhan Zuo, Ruoxuan Zhang, and Wen-Huang Cheng. Emoart: A multidimensional dataset for emotion-aware artistic generation. In *ACM MM*, pages 12644–12650, 2025. 3
- [47] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, pages 3836–3847, 2023. 2, 4
- [48] Peixuan Zhang, Shuchen Weng, Jiajun Tang, Si Li, and Boxin Shi. Towards deeper emotional reflection: Crafting affective image filters with generative priors. *IEEE transactions on pattern analysis and machine intelligence*, 2025. 3
- [49] Peixuan Zhang, Shuchen Weng, Chengxuan Zhu, Binghao Tang, Zijian Jia, Si Li, and Boxin Shi. Affective image editing: Shaping emotional factors via text descriptions. *International Journal of Computer Vision*, 134(1):16, 2026. 3
- [50] Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. Inversion-based style transfer with diffusion models. In *CVPR*, pages 10146–10156, 2023. 3, 6, 7
- [51] Zhanjie Zhang, Quanwei Zhang, Huaizhong Lin, et al. Towards highly realistic artistic style transfer via stable diffusion with step-aware and layer-aware prompt. *arXiv preprint arXiv:2404.11474*, 2024. 6
- [52] Chenxi Zhao, Jinglei Shi, Liqiang Nie, and Jufeng Yang. To err like human: Affective bias-inspired measures for visual emotion recognition evaluation. In *NeurIPS*, pages 134747–134769, 2024. 2
- [53] Sicheng Zhao, Xiaopeng Hong, Jufeng Yang, Yanyan Zhao, and Guiguang Ding. Toward label-efficient emotion and sentiment analysis. *IEEE*, 111(10):1159–1197, 2023. 2