

# Towards Deeper Emotional Reflection: Crafting Affective Image Filters with Generative Priors

Peixuan Zhang<sup>†</sup>, Shuchen Weng<sup>†</sup>, Jiajun Tang, Si Li\*, and Boxin Shi, *Senior Member, IEEE*

**Abstract**—Social media platforms enable users to express emotions by posting text with accompanying images. In this paper, we propose the Affective Image Filter (AIF) task, which aims to reflect visually-abstract emotions from text into visually-concrete images, thereby creating emotionally compelling results. We first introduce the AIF dataset and the formulation of the AIF models. Then, we present AIF-B as an initial attempt based on a multi-modal transformer architecture. After that, we propose AIF-D as an extension of AIF-B towards deeper emotional reflection, effectively leveraging generative priors from pre-trained large-scale diffusion models. Quantitative and qualitative experiments demonstrate that AIF models achieve superior performance for both content consistency and emotional fidelity compared to state-of-the-art methods. Extensive user study experiments demonstrate that AIF models are significantly more effective at evoking specific emotions. Based on the presented results, we comprehensively discuss the value and potential of AIF models.

**Index Terms**—Emotion analysis, Image editing, Image filters.

## I. INTRODUCTION

SOCIAL media platforms (e.g., Facebook and X), as widely embraced communication tools, enable people to forge emotional connections with one another [1]. By sending posts, users can share what they are seeing, talking, and thinking. Although texts in posts can express emotions straightforwardly, attached images may act as more powerful affective stimuli to attract the attention of their followers.

In an effort to help create unique and emotionally compelling images that stand out from the crowd, the task of Affective Image Filter (AIF) [2] has emerged as an innovative research topic. As shown in Fig. 1, given a content image that depicts users’ daily experience and a text description that reflects personal thoughts and feelings, an AIF model can reflect visually-abstract emotions from user-provided texts to visually-concrete images with appropriate colors and textures. Specifically, a qualified AIF model should exhibit both *content consistency* and *emotional fidelity*, which requires the model to effectively preserve the details of content images, accurately understand emotions in text descriptions, and appropriately visualize colors and textures with corresponding emotions.

P. Zhang and S. Li are with the School of Artificial Intelligence, Beijing University of Posts and Telecommunications, Beijing 100876, China. Email: {pxzhang, lisi}@bupt.edu.cn.

S. Weng is with the Beijing Academy of Artificial Intelligence, Beijing 100083, China. Email: {scweng}@baai.ac.cn.

J. Tang and B. Shi are with the State Key Laboratory of Multimedia Information Processing and National Engineering Research Center of Visual Technology, School of Computer Science, Peking University, Beijing 100871, China. Email: {jiajun.tang, shiboxin}@pku.edu.cn.

<sup>†</sup> Equal contribution. \* Corresponding author.

The previous Basic AIF model [2] (denoted as AIF-B) achieves the aforementioned objectives with a specifically collected AIF dataset. Since low-level and mid-level features (e.g., colors and styles) can effectively represent evoked emotions [3], [4], the AIF dataset focuses on famous paintings across various aesthetic styles. Based on it, AIF-B takes the first step to concrete visualization of emotions, formulating the process as three parts: (i) **Model architecture**. AIF-B is based on a multi-modal transformer architecture, which unifies content images and text descriptions into the token representation, enabling the model to learn their interdependency through subsequent transformer blocks. (ii) **Emotional understanding**. To understand inherent emotional properties, AIF-B leverages the prior knowledge of the VAD dictionary [5], which encodes the valence, arousal, and dominance scores for emotional words. Additionally, AIF-B introduces an emotional distribution loss to capture high-dimensional emotional cues. (iii) **Emotional visualization**. To visualize the corresponding emotional expressions in synthesized images, AIF-B introduces a sentiment metric loss and an anchor-based sentiment loss to effectively reflect specific emotions. Content loss, style loss, GAN loss, and identity loss are further adopted to improve the aesthetic quality of synthesized images.

Although AIF-B has made remarkable progress in applying the corresponding filters, three major challenges still remain for this initial approach: (i) **Blurry details**. When applying AIF-B, a potential side effect is the removal of high-frequency details from objects, leading to a blurry effect (Fig. 1(a), synthesizing windows with blurred contours). (ii) **Keyword blindness**. To accurately evoke emotions, AIF-B may overly focus on specific keywords within text descriptions, potentially missing the broader context (Fig. 1(b), expressing feelings with less warmth and relaxation). (iii) **Overly-stylized effects**. When reflecting emotions to images, AIF-B may excessively apply artistic modifications, resulting in abrupt colors, exaggerated contrast, and distorted content (Fig. 1(c), visualizing emotions with over-saturated colors).

In this paper, we extend our preliminary work AIF-B [2] towards deeper emotional reflection. We design the Deeper AIF model (denoted as AIF-D) following the aforementioned formulation of AIF-B: (i) **Refined model architecture**. We build AIF-D upon the pre-trained large-scale generative model [6], leveraging its rich generative priors to produce high-quality images. To address the blurry details caused by the downsampling block during image encoding, we design a content preservation module that extracts multi-scale spatial

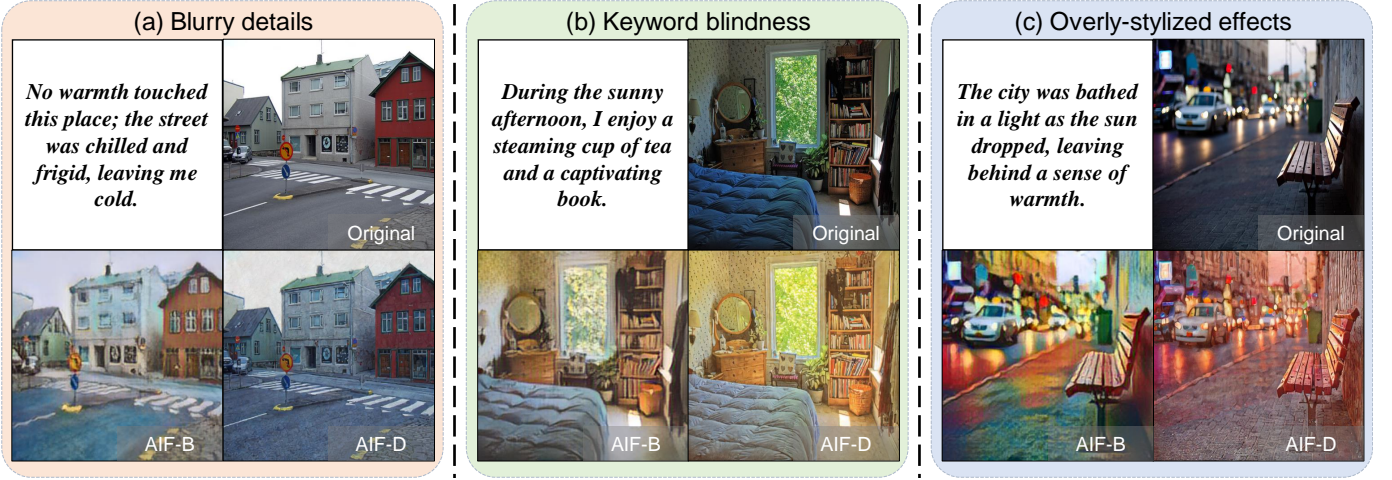


Fig. 1. Towards deeper emotional reflection, we propose AIF-D to better meet the requirements of emotional fidelity and content consistency. Compared to previous AIF-B [2], AIF-D overcomes three major challenges: (a) **Blurry details**. AIF-D preserves high-frequency details of the content image more effectively. (b) **Keyword blindness**. AIF-D has a deeper emotional understanding of text descriptions that lack specific keywords. (c) **Overly-stylized effects**. AIF-D reflects emotions with more appropriate artistic representation.

features. This module effectively preserves these details and maintains content consistency (Fig. 1 (a), the intricate details of the houses). (ii) **Nuanced emotional understanding**. To avoid keyword blindness when handling complex emotional expressions, we leverage a Large Language Model (LLM) [7], [8] and introduce Chain-of-Thought (CoT) prompting [9], [10] to better reflect emotions from text descriptions. Additionally, we propose a voting ensemble mechanism that evaluates emotional distributions from multiple perspectives. By ensembling estimated emotional distributions, this mechanism enables accurate analysis of evoked emotions (Fig. 1 (b), cozy atmosphere for the room). (iii) **Advanced emotional visualization**. To enable accurate visualization of emotions, we propose an emotional reflection strategy to effectively learn both emotional and artistic priors and present compelling artistic features. After that, we redesign the aesthetic loss to balance artistic style and content consistency, ensuring that the synthesized images align more closely with the objectives of the AIF task. (Fig. 1 (c), appropriate artistic representation in the street view).

In summary, AIF-D improves content consistency and emotional fidelity, advancing the AIF models towards deeper emotional reflection by making the following contributions:

- We leverage the generative priors from the pre-trained generative model and propose a content preservation module to preserve high-frequency details.
- We integrate an LLM and CoT prompting to improve emotional understanding. The voting ensemble mechanism is presented to accurately analyze evoked emotions.
- We propose an emotional reflection strategy to enhance the image decoder. The aesthetic loss is further redesigned to balance artistic style and content consistency.

The remainder of this paper is organized as follows. We begin with an introduction to the related work in Sec. II, including text-guided image editing, visual-emotional analysis, and image style transfer. We introduce the AIF dataset and AIF model formulation in Sec. III. We then introduce AIF-B in Sec. IV,

followed by AIF-D in Sec. V, highlighting improvements made in model architecture, emotional understanding, and emotion visualization. In Sec. VI, we conduct both quantitative and qualitative experiments. In Sec. VII, we explore applications, failure cases, robustness testing, and differences from general editing methods. Finally, the paper is concluded in Sec. VIII.

## II. RELATED WORK

### A. Text-guided Image Editing.

As a user-friendly option, text descriptions have been widely used in image editing tasks, *e.g.*, enabling modifications across color [11], [12], texture [2], [13], and high-level semantic [14], [15]. Initial research efforts are directed towards developing effective text injection modules [16]–[18]. The emergence of language-image pre-trained models [19] further expands the scope of image editing, allowing for a broader, open-vocabulary approach [20]–[22]. Recently, diffusion models [23], [24] have demonstrated remarkable improvement in text-to-image generation task [6], [25], [26], inspiring researchers to leverage the powerful generative capabilities to advanced image editing applications, *e.g.*, a fine-tuning model for multi-conditional image editing [27], [28], personalizing generative models through image inversion [29], [30], and developing advanced sampling strategies to steer generation results [31], [32]. Leveraging rich generative priors from pre-trained models, AIF models could present rich details and create visually compelling synthesizing results.

### B. Visual Emotional Analysis.

The visual emotional analysis holds a crucial position in computer vision [33]. In the early years, relevant efforts [3] focus on utilizing hand-crafted features for emotional recognition in images. Various low-level features are adopted by numerous studies [4], [34] to represent emotional cues, *e.g.*, color, shape, and texture. The potential of mid-level features in emotional recognition is subsequently identified [35]–[37],

providing a more intuitive human understanding than low-level features. The emergence of deep learning has revolutionized this domain, researchers [38]–[41] extract high-level features to improve performance in image-based emotional recognition by focusing on content-rich information, *e.g.*, facial expressions. Following them, multiple emotional features are integrated for a comprehensive understanding [42], [43]. Furthermore, recent works [44], [45] introduce contextual understanding for visual emotion analysis, since emotions are conveyed through scene interplay, social interactions, and the overall environment. Recent large language models (LLMs) have led to notable advancements in understanding text descriptions and generating corresponding responses [7], [8]. With Chain-of-Thought (CoT) prompting [9], [10], LLMs are improved towards deeper emotional understanding. Inspired by these pioneering works, AIF models ensemble low-level and mid-level features as cues, leveraging emotional reasoning to accurately analyze nuanced emotions within text descriptions. High-level features are omitted since an image filter should preserve the fundamental content or narrative of images.

### C. Image Style Transfer.

Early style transfer models focus on replicating individual styles through optimization-based methods [46] or designing end-to-end convolutional neural networks [47], [48]. Later, arbitrary style transfer models [49]–[52] adaptively synthesize stylized results from any given reference image. To improve the ability of contextual understanding, attention mechanisms are integrated [53]–[56], establishing long-range dependencies between image patches. In recent years, diffusion models [6], [27] break new ground in producing high-fidelity images, which facilitates style transfer research to incorporate the generative capabilities of these models [57]–[59]. Meanwhile, text descriptions are gradually replacing traditional reference images, offering a more flexible and user-friendly approach to style transfer [13], [22], [60]. Compared to style transfer methods that rely on explicit color and texture features, an AIF model should reflect implicit emotional cues.

## III. COMMON SETUPS FOR AIF TASK

The AIF task enables users to reflect their thoughts and feelings, which requires the model to understand visually-abstract emotions from user-provided text descriptions and reflect them into visually-concrete images with appropriate colors and textures. We introduce the AIF dataset to provide comprehensive emotional data. Additionally, we formulate the AIF model with four key components.

### A. AIF Dataset Collection

Although previous datasets (*e.g.*, ArtEmis [61] and ArtEmis v2 [62]) connect visual elements, text descriptions, and emotions, their descriptions may emphasize the visual content over the evoked emotions of artistic images. This limitation motivates us to create the AIF dataset to prepare qualified data for training and evaluating relevant models. Specifically, we first merge all images from ArtEmis [61] and ArtEmis v2 [62]. Then,

we manually filter the unqualified text descriptions, remaining those that illustrate evoked emotions and discarding those that focus on visual details. After that, we ensure each image in the AIF dataset has an average of four to five such emotional descriptions, and categorize these descriptions according to Mikel’s wheel [63]. Due to the inherent subjectivity and ambiguity of emotions, we represent the evoked emotions of each image as a distribution across these categories, enabling the model to capture the relationships between emotions.

Finally, the AIF dataset consists of abundant aesthetic samples as anchor images that have diverse colors and textures, used as fixed points of visually-concrete references for visualizing colors and textures from visually-abstract emotions. Each image is paired with several corresponding text descriptions associated with the emotional category labels in Mikel’s wheel [63]. We directly adopt the original emotional categories provided by the source datasets. To ensure the quality and reliability of these annotations, we conduct a second round of verification on a 10% subset of our dataset. After measuring the inter-rater agreement using Cohen’s Kappa [64], the result is 0.81, indicating almost perfect agreement and confirming the reliability of the annotations. We split the dataset into a training set (69.6K images and 292.9K text descriptions) and an evaluation set (7.7K images and 32.5K text descriptions). To evaluate the performance of relevant methods, we use samples from the COCO-Stuff [65] dataset as content images.

### B. AIF Model Formulation

We systematically formulate an AIF model with the following four key components:

1) *Inputs*: Content images  $I^{\text{cnt}} \in \mathbb{R}^{H \times W \times 3}$  to provide visual content and text descriptions  $T^{\text{des}} \in \mathbb{N}^M$  to evoke specific emotions, where  $H$  and  $W$  are the image size, and  $M$  is the number of words in the text descriptions.

2) *Crafting architecture with content preservation*: Content images  $I^{\text{cnt}}$  and text descriptions  $T^{\text{des}}$  are encoded by an image encoder  $\mathcal{E}$  and a text encoder  $\mathcal{T}$ , respectively:

$$Z^{\text{img}} = \mathcal{E}(I^{\text{cnt}}), \quad Z^{\text{tex}} = \mathcal{T}(T^{\text{des}}), \quad (1)$$

where  $Z^{\text{img}} \in \mathbb{R}^{h \times w \times C_{\text{img}}}$  and  $Z^{\text{tex}} \in \mathbb{R}^{M \times C_{\text{tex}}}$  are separately image tokens and text tokens.  $h = H/f$ ,  $w = W/f$ , and  $f$  is a hyper-parameter of the downsampling rate.  $C_{\text{img}}$  and  $C_{\text{tex}}$  are the numbers of embedding channels of image and text tokens, respectively. To capture long-range dependencies between tokens across different modalities, both the image tokens  $Z^{\text{img}}$  and text tokens  $Z^{\text{tex}}$  are fed into a generation model  $G$  to preserve the overall structure of content images:

$$Z' = G(Z^{\text{img}}, Z^{\text{tex}}). \quad (2)$$

An image decoder  $\mathcal{D}$  is further used to synthesize images  $I^{\text{out}} \in \mathbb{R}^{H \times W \times 3}$  that effectively reflect specific emotional responses from human observers:

$$I^{\text{out}} = \mathcal{D}(Z'). \quad (3)$$



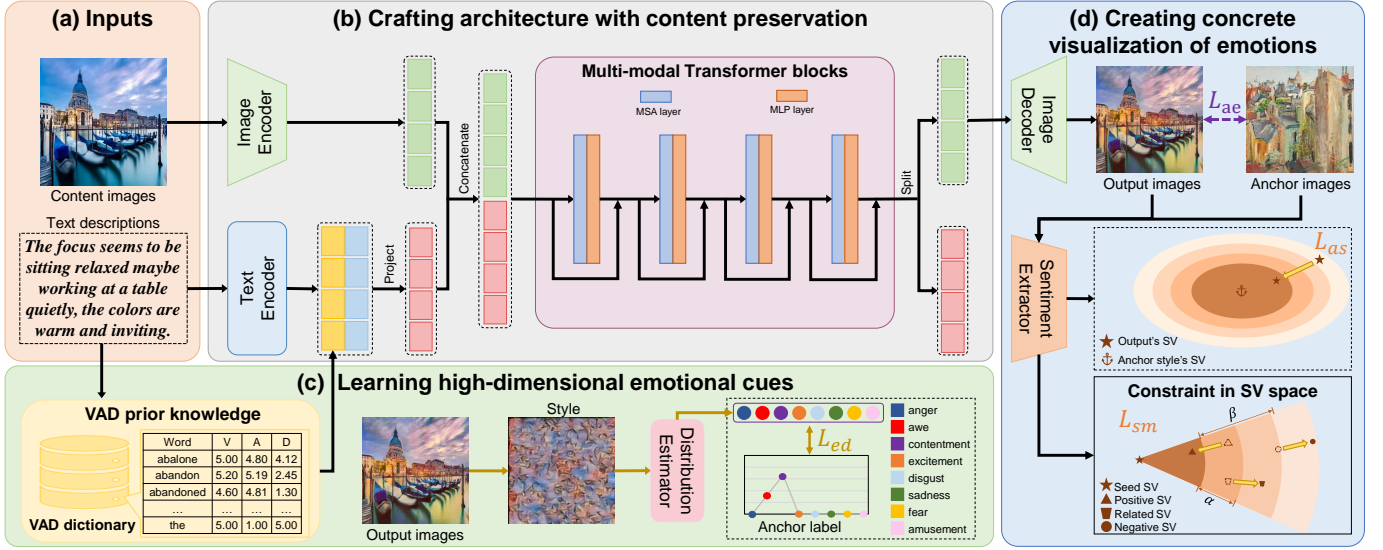


Fig. 2. The pipeline of AIF-B [2]. (a) Users input content images to provide visual content and text descriptions to evoke specific emotions. (b) Content images and text descriptions are encoded by the image encoder and text encoder, respectively. The encoded image and text tokens are fed into multi-modal transformer blocks, each comprising a Multi-headed Self-Attention (MSA) layer, a Multi-Layer Perceptron (MLP) layer, and residual connections (Sec. IV-A). (c) To understand inherent emotional properties, AIF-B leverages the VAD dictionary as emotional prior knowledge. The emotional distribution loss is introduced to capture high-dimensional emotional cues (Sec. IV-B). (d) The sentiment metric loss and anchor-based sentiment loss are proposed to effectively reflect specific emotions (Sec. IV-C). The aesthetic loss is further adopted to ensure the aesthetic quality and enrich the artistic style of synthesized images (Sec. IV-D).

3) *Learning high-dimensional emotional cues*: To understand inherent emotional properties, we augment the emotional expression of text tokens  $Z^{\text{emo}}$  by incorporating external emotional cues:

$$Z^{\text{emo}} = P(Z^{\text{tex}}), \quad (4)$$

where  $P$  represents an emotional augmentation module. We further replace  $Z^{\text{tex}}$  with  $Z^{\text{emo}}$  in Eq. (2) to leverage these emotional cues. Since emotions are ambiguous and subjective, an emotional distribution loss  $\mathcal{L}_{\text{ed}}$  is required to capture the synthesized emotional cues, which ensures the evoked emotional distributions align with users' intentions.

4) *Creating concrete emotional visualization*: The AIF model is equipped with four losses to synthesize visually appealing images. To ensure the synthesized images could evoke specific emotional responses, a sentiment metric loss  $\mathcal{L}_{\text{sm}}$  is required to learn relationships between emotions. To create appropriate colors and textures, an anchor-based sentiment loss  $\mathcal{L}_{\text{as}}$  is required to learn artistic representations from corresponding anchor images. Additionally, to improve the aesthetic quality of synthesized images, the aesthetic loss  $\mathcal{L}_{\text{ae}}$  is required to balance artistic style with content consistency.

#### IV. BASIC AIF MODEL

In this section, AIF-B is introduced as an initial attempt based on the AIF model formulation (Sec. III-B). We first present the multi-modal transformer architecture (Sec. IV-A). Then, we fetch the emotional cues from the VAD dictionary and design the emotional distribution loss (Sec. IV-B). Next, we introduce the sentiment loss and the anchor-based sentiment loss (Sec. IV-C), along with the aesthetic loss (Sec. IV-D), for effective emotional visualization. Finally, we provide the details of training settings (Sec. IV-E).

##### A. Multi-modal Transformer

As illustrated in Fig. 2 (b), given sampled batches, the multi-modal transformer separately encodes text descriptions and content images, and captures their dependencies in the subsequent transformer blocks. Specifically, the multi-modal transformer can be divided into four components.

**Image encoder.** Content images  $I^{\text{cnt}}$  are split into image patches and a transformer-based image encoder is adopted to extract image tokens  $Z^{\text{img}} = [Z_1^{\text{img}}, \dots, Z_N^{\text{img}}] \in \mathbb{R}^{N \times C_{\text{img}}}$ , where  $N$  is the number of image patches and  $C_{\text{img}}$  is the number of embedding channels.

**Text encoder.** Text descriptions are encoded with pre-trained BERT [66] to form text tokens  $Z^{\text{tex}} = [Z_1^{\text{tex}}, \dots, Z_M^{\text{tex}}] \in \mathbb{R}^{M \times C_{\text{tex}}}$ , where  $M$  is the number of words and  $C_{\text{tex}}$  is the number of embedding channels.

**Multi-modal transformer.** A stack of MLP layers further projects the image and text tokens into a shared latent space. Then, modal-type embeddings  $Z_0^{\text{typ}}$  and  $Z_1^{\text{typ}}$  are introduced to distinguish token modalities:

$$\hat{Z}^{\text{img}} = \text{MLP}(Z^{\text{img}}) + Z_0^{\text{typ}}, \quad \hat{Z}^{\text{tex}} = \text{MLP}(Z^{\text{tex}}) + Z_1^{\text{typ}}, \quad (5)$$

These embeddings are then concatenated as the initial input of the transformer:

$$Z_0 = [\hat{Z}_1^{\text{img}}, \dots, \hat{Z}_N^{\text{img}}, \hat{Z}_1^{\text{tex}}, \dots, \hat{Z}_M^{\text{tex}}] \in \mathbb{R}^{(N+M) \times C_0}, \quad (6)$$

where  $C_0$  is the number of channel. The adopted transformer consists of  $L$  transformer blocks [67], and each block includes a self-attention layer, an MLP layer, and two residual connections. The computation within each block is formulated as:

$$[\bar{Z}_i] = \text{MSA}(\text{LN}([Z_{i-1}])) + [Z_{i-1}], \quad i \in \{1, \dots, L\} \quad (7)$$

$$[Z_i] = \text{MLP}(\text{LN}([\bar{Z}_i])) + [\bar{Z}_i], \quad i \in \{1, \dots, L\} \quad (8)$$



where LN means the LayerNorm.

**Image decoder.** A three-layer CNN is used as the image decoder to synthesize results  $I^{\text{out}}$ , which alleviates the grid artifacts caused by the patch partition of images.

### B. External Emotional Cues

As illustrated in Fig. 2 (c), the VAD dictionary, as an emotional augmentation module, augments emotional expression. The emotional distribution loss is employed to evoke emotions accurately.

**VAD prior knowledge.** Affective tokens are fetched from the VAD dictionary [5] and used as the emotional cues to extract the inherent emotional properties of each word (*i.e.*, valence, arousal, and dominance). Specifically, valence indicates the positivity or negativity of the emotion, arousal reflects the intensity or excitement level of the emotion, and dominance represents the degree of control or influence one feels with the emotion. Words not included in the VAD dictionary are assigned neutral values to ensure the completeness and consistency of emotional understanding. These tokens are denoted as  $Z^{\text{vad}} \in \mathbb{R}^{M \times C_{\text{vad}}}$ , where  $C_{\text{vad}}$  is the number of embedding channels. To augment the emotional expression, text tokens are concatenated with affective tokens as  $Z^{\text{emo}} = [Z^{\text{tex}}; Z^{\text{vad}}] \in \mathbb{R}^{M \times (C_{\text{tex}} + C_{\text{vad}})}$ , replacing  $Z^{\text{tex}}$  in Eq. (5) and Eq. (6).

**Emotional distribution loss.** To effectively capture synthesized emotional cues, an emotional distribution loss is proposed. Since emotions are inherently ambiguous and subjective, a single image can evoke a diverse range of emotional responses even within the same individual. Instead of categorizing a single dominant emotion for each image, the emotional distribution loss aims to estimate the comprehensive emotional distribution. Specifically, a VGG-based distribution estimator  $\varphi$  is used to extract color and style features through the mean and variance of the feature representations. The emotional distribution loss is defined using the Kullback-Leibler (KL) divergence [68], which minimizes the discrepancy between the emotional distributions of synthesized images and the ground truth:

$$\mathcal{L}_{\text{ed}} = \sum_{i=1}^{N_{\text{cat}}} d_i \ln \frac{d_i}{\varphi(I^{\text{out}})_i}, \quad (9)$$

where  $N_{\text{cat}}$  is the number of emotional categories,  $\varphi(I^{\text{out}})_i$  and  $d_i$  are the value of  $i$ -th category of the estimated distribution and the ground truth, respectively.

### C. Emotional Reflection Loss

As demonstrated in Fig. 2 (d), the sentiment metric loss  $\mathcal{L}_{\text{sm}}$  and anchor-based sentiment loss  $\mathcal{L}_{\text{as}}$  are proposed to enable synthesized images to effectively reflect specific emotions.

**Sentiment metric loss.** AIF-B learns relationships between emotions in three steps: (i) Following Yang [69], a sentiment extractor is built. A VGG network is used to extract multi-level features of synthesized images. These features are then projected through a convolutional layer and used to calculate multiple Gram matrices at multiple feature levels. The sentiment vector  $V$  is then defined by concatenating selected elements from these Gram matrices:

$$V = \text{Concat}_{i \in \{1, \dots, N_{\text{gram}}\} \cup j \in \{1, \dots, N_{\text{lev}}\}} (\Phi_{i,j}), \quad (10)$$

where  $\Phi_{i,j}$  means the  $i$ -th upper triangular elements in the Gram matrix at  $j$ -th feature level,  $N_{\text{gram}}$  is the number of elements, and  $N_{\text{lev}}$  is the number of levels. (ii) Based on Mikel's wheel [63] that defines relationship between emotions, and given randomly sampled text descriptions  $T^{\text{sed}}$ , tuples of text samples  $[T^{\text{sed}}, T^{\text{pos}}, T^{\text{rel}}, T^{\text{neg}}]$  are constructed. Here,  $T^{\text{pos}}$  represents positive samples from the same emotional region,  $T^{\text{rel}}$  represents related samples from adjacent emotional regions, and  $T^{\text{neg}}$  represents negative samples from the opposite emotional region. The corresponding sentiment vectors are denoted as  $[V^{\text{sed}}, V^{\text{pos}}, V^{\text{rel}}, V^{\text{neg}}]$ . (iii) Defining the emotional distance  $F_{\text{dis}}$  as the minimum number of steps between emotion regions in the Mikel's wheel [63], the distance between sentiment vectors is formulated as:  $F_{\text{sw}}(V_i, V_j) = \frac{\|V_i - V_j\|^2}{F_{\text{dis}}(V_i, V_j)}$ . As a result, the sentiment metric loss is formulated in a metric-learning manner:

$$\mathcal{L}_{\text{sm}} = \max(F_{\text{sw}}(V^{\text{sed}}, V^{\text{pos}}) - F_{\text{sw}}(V^{\text{sed}}, V^{\text{rel}}) + \alpha, 0) + \max(F_{\text{sw}}(V^{\text{sed}}, V^{\text{rel}}) - F_{\text{sw}}(V^{\text{sed}}, V^{\text{neg}}) + \beta, 0), \quad (11)$$

where  $\alpha = 0.02$  and  $\beta = 0.01$  are hyper-parameters that control margins between sentiment vectors.

**Anchor-based sentiment loss.** To ensure that synthesized images have appropriate colors and textures, the sentiment extractor is further used to obtain sentiment vectors for each synthesized images and their corresponding anchor images (denoted as  $V^{\text{out}}$  and  $V^{\text{acr}}$ ). By minimizing the mean squared error between these vectors, the model aligns the synthesized images with the desired artistic representations:

$$\mathcal{L}_{\text{as}} = \|V^{\text{out}} - V^{\text{acr}}\|_2. \quad (12)$$

### D. Visual Aesthetics Loss

As depicted in Fig. 2 (d), the aesthetic loss  $\mathcal{L}_{\text{ae}}$  is designed to improve the aesthetic quality of synthesized images. In practice, AIF-B defines aesthetic loss as a composite function, including content loss  $\mathcal{L}_c$  [49], style loss  $\mathcal{L}_s$  [53], GAN loss  $\mathcal{L}_{\text{GAN}}$  [70], and identity loss  $\mathcal{L}_{\text{id}}$  [71], detailed as follows:

**Content loss.** To preserve the visual features and overall structure, the content loss uses a pre-defined VGG network [72] to extract multi-level features and narrows the squared error between corresponding features at each level as:

$$\mathcal{L}_c = \sum_i \|\phi_i^{\text{out}} - \phi_i^{\text{cnt}}\|_2, \quad (13)$$

where  $\phi_i^{\text{out}}$  and  $\phi_i^{\text{cnt}}$  are extracted features of synthesized images and content images at  $i$ -th level, respectively.

**Style loss.** Following the style transfer methods [49], [53], we represent the image style using the mean and variance of extracted feature maps. To transfer the artistic style of images, the style loss minimizes the style difference between feature statistics of synthesized images and anchor images at each level as:

$$\mathcal{L}_s = \sum_i \|\mu(\phi_i^{\text{out}}) - \mu(\phi_i^{\text{acr}})\|_2 + \|\sigma(\phi_i^{\text{out}}) - \sigma(\phi_i^{\text{acr}})\|_2, \quad (14)$$

where  $\mu$  and  $\sigma$  denote the mean and variance functions, respectively.  $\phi_i^{\text{acr}}$  is extracted features of anchor images.

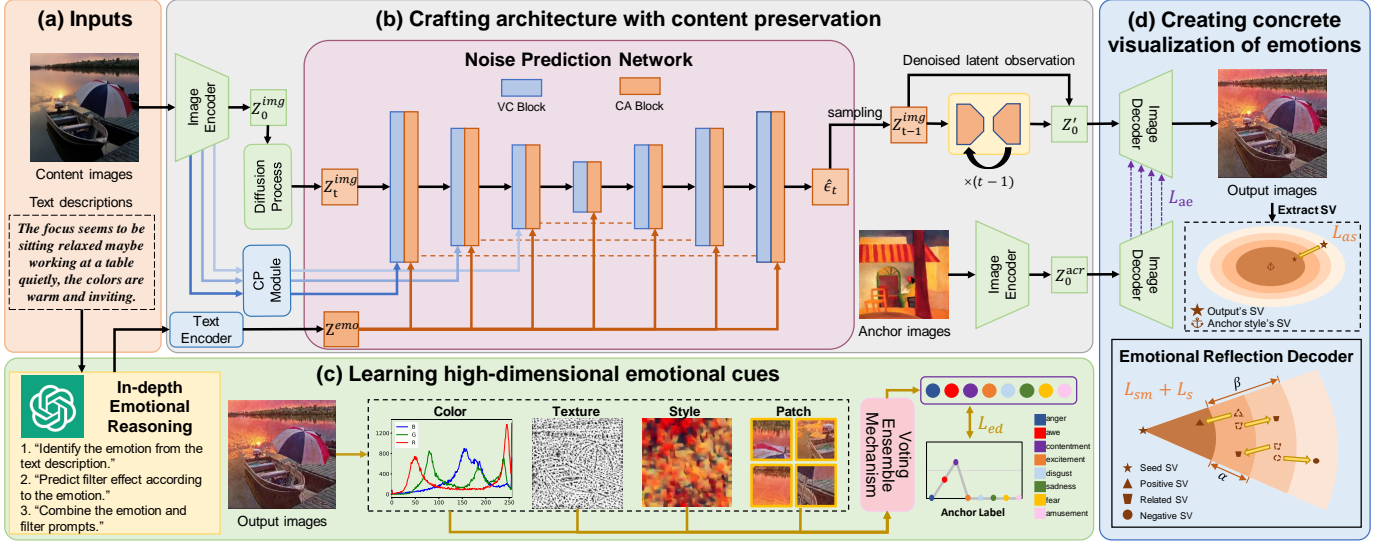


Fig. 3. The pipeline of AIF-D. (a) Users input content images to provide visual content and text descriptions to evoke specific emotions. (b) Content images and text descriptions are fed into the image encoder and text encoder to extract image tokens and text tokens, respectively. The noise prediction network estimates the noise at each diffusion step. Within the downsampling modules, each Vanilla Convolution (VC) block before the Cross-Attention (CA) blocks is equipped with a Content Preservation (CP) module. These CP modules integrate the context of content images to preserve high-frequency details (Sec. V-A). (c) An LLM and CoT prompting are used to handle complex emotional expressions with in-depth emotional reasoning. The voting ensemble mechanism is introduced to evaluate emotional distributions from different points of view and ensembles low-level (colors and texture) and mid-level features (image style and patch features) to enable accurate analysis of evoked emotions (Sec. V-B). (d) To enable accurate visualization of emotions, we propose an emotional reflection strategy to enhance the emotional understanding of the image decoder, using the sentiment metric loss and anchor-based sentiment loss (Sec. V-C). Finally, images are synthesized using a redesigned aesthetic loss, achieving a balance between artistic style and content consistency (Sec. V-D).

**GAN loss.** To ensure consistency between synthesized images and text descriptions while maintaining aesthetic quality, a conditional-unconditional discriminator is introduced:

$$\mathcal{L}_{GAN} = \log D(I^{acr}) + \log (1 - D(G(I^{cnt}, Z^{emo}))) + \log D(I^{acr}, Z^{emo}) + \log (1 - D(G(I^{cnt}, Z^{emo}), Z^{emo})), \quad (15)$$

where  $D$  and  $G$  are the discriminator and generator.

**Identity loss.** To learn richer semantic representations (e.g., color, texture, and visual content), anchor images and corresponding text descriptions are fed into the AIF model to synthesize identity images. The identity loss ensures consistency between identity images and their original anchor images:

$$\mathcal{L}_{id} = \|I^{idt} - I^{acr}\|_2 + \lambda \sum_i \|\phi_i^{idt} - \phi_i^{acr}\|_2, \quad (16)$$

where  $I^{idt}$  and  $\phi_i^{idt}$  are identity images and extracted features at  $i$ -th level, respectively.  $I^{acr}$  means corresponding anchor images.  $\lambda = 0.01$  is a hyper-parameter.

Consequently, the aesthetic loss in AIF-B is defined as:

$$\mathcal{L}_{ae} = \lambda_c \mathcal{L}_c + \lambda_s \mathcal{L}_s + \lambda_{gan} \mathcal{L}_{GAN} + \lambda_{id} \mathcal{L}_{id}, \quad (17)$$

where  $\lambda_c = 5$ ,  $\lambda_s = 0.3$ ,  $\lambda_{gan} = 3$ , and  $\lambda_{id} = 2$  denote the weights for losses.

#### E. Learning and training details

**Losses.** AIF-B is trained by solving a minimax optimization problem involving full objective losses:

$$\max_D \min_G \lambda_{ed} \mathcal{L}_{ed} + \lambda_{sm} \mathcal{L}_{sm} + \lambda_{as} \mathcal{L}_{as} + \lambda_{ae} \mathcal{L}_{ae},$$

where  $\lambda_{ed} = 140$ ,  $\lambda_{sm} = 30$ ,  $\lambda_{as} = 600$ , and  $\lambda_{ae} = 1$  are hyper-parameters, which are determined empirically.

**Training Details.** AIF-B is trained on 4 NVIDIA TITAN RTX GPUs for 80K iterations with batch size 24 for 30 hours. We use the Adam optimizer minimizes losses with the warm-up adjustment strategy [73].

## V. DEEPER AIF MODEL

Although AIF-B takes the first step to concrete visualization of emotions, it still suffers from blurry details, keyword blindness, and overly-stylized effects. In this section, we present AIF-D as an extension of AIF-B, advancing towards deeper emotional reflection. For clarity, we follow the same presentation structure as AIF-B (Sec. IV) to describe AIF-D. We begin by introducing the diffusion-model backbone with a content preservation module (Sec. V-A). Then, we leverage an LLM for in-depth emotional reasoning while ensembling emotional cues from different perspectives (Sec. V-B). We improve the emotional understanding of the image decoder based on the sentiment metric and anchor-based sentiment loss (Sec. V-C). The aesthetic loss is further proposed to balance artistic style and content consistency (Sec. V-D). Finally, we provide the details of training settings (Sec. V-E).

#### A. Diffusion-model backbone

As shown in Fig. 3 (b), AIF-D is built upon the pre-trained latent diffusion model [6] to leverage its rich generative priors, rather than training a multi-modal transformer from scratch (Sec. IV-A). The model can be divided into four components. **Condition encoders.** A pre-trained Variational Autoencoder (VAE) [74] is adopted as the image encoder  $\mathcal{E}$ , converting content images  $I^{cnt}$  into image tokens  $Z^{img}$ . Text descriptions  $T^{des}$

are processed through the CLIP [19], as the text encoder  $\mathcal{T}$ , to generate text tokens  $Z^{\text{tex}}$ . These tokens are then injected into the latent space of a UNet-based noise prediction network  $\epsilon_\theta$ , which incorporates the cross-attention mechanism [75].

**Content preservation module.** As the condition encoder compresses the image, high-frequency details are mostly lost, resulting in blurry details and reduced content consistency. To address this degradation, the integration of a content preservation module is proposed within the noise prediction network, leveraging intermediate encoder features as the additional context of the content image. Specifically, this module injects multi-scale features from the image encoder into a modified U-Net architecture [6], serving as contextual anchors that guide the noise prediction. By injecting these intermediate features into corresponding stages of the U-Net backbone, the content preservation module preserves high-frequency details and improves content consistency. This process could be expressed as:

$$f'_i = f_i + F_{\text{cp}}(\mathcal{E}_i), \quad i \in \{1, \dots, N_{\text{enc}}\}, \quad (18)$$

where  $f_i$  and  $\mathcal{E}_i$  are the features in the U-Net backbone and the image encoder of the  $i$ -block respectively,  $F_{\text{cp}}$  means a stack of downsampling convolution, and  $N_{\text{enc}}$  is the number of blocks in the image encoder.

**Forward process.** Given the property of the forward process in diffusion models [23], [24], the forward process plays a critical role in transforming the image tokens into a latent space, learning a robust generative pathway. Specifically, let  $Z_0^{\text{img}}$  denote the encoded image features  $Z^{\text{img}}$  at timestep  $t = 0$ . The noised sample  $Z_t^{\text{img}}$  can be expressed as a linear combination of  $Z_0^{\text{img}}$  and a noise variable  $\epsilon$  as the diffusion process:

$$Z_t^{\text{img}} = \sqrt{\alpha_t} Z_0^{\text{img}} + \sqrt{1 - \alpha_t} \epsilon, \quad (19)$$

where  $\epsilon \sim \mathcal{N}(0, 1)$  is the Gaussian noise and  $\alpha_t$  is the parameter for controlling noise levels.

**Backward process.** The noise prediction network aims to iteratively refine the noised sample  $Z_t^{\text{img}}$  until it converges to the clear content latent code  $Z_0'$ :

$$Z_{t-1}^{\text{img}} = \frac{1}{\sqrt{\alpha_t}} \left( Z_t^{\text{img}} - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(Z_t^{\text{img}}, t, Z^{\text{tex}}, Z_0^{\text{img}}) \right) + \sigma_t \epsilon, \quad (20)$$

where  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$  controls the amount of noise,  $Z^{\text{tex}}$  denotes text tokens,  $\sigma_t$  defines the standard deviation of the noise. Finally, a pre-trained VAE image decoder  $\mathcal{D}$  is used to map the latent code back into the pixel space and present the synthesized image as  $I^{\text{out}} = \mathcal{D}(Z_0')$ .

### B. Ensembling Emotional Cues

As illustrated in Fig. 3 (c), AIF-D replaces the VAD dictionary (Sec. IV-B) with an LLM and CoT prompting as the emotional augmentation module to facilitate in-depth emotional reasoning. Instead of estimating emotional distribution solely from style (Sec. IV-B), the voting ensemble mechanism extracts image features from different perspectives to ensure the synthesized results evoke the corresponding emotions.

**In-depth emotional reasoning.** Relying solely on individual emotional words can lead to keyword blindness and failure

to capture complex emotional expressions. Towards deeper understanding of emotions, an LLM [8] and the CoT prompting [9], [10] are introduced to facilitate in-depth emotional reasoning. Specifically, the CoT prompting process enhances user-provided text descriptions into a rich emotional prompt through three sequential steps: (i) First, we prompt the LLM [8] to analyze the text description and extract emotional keywords (e.g., calmness, contentment, and enjoyment). (ii) Next, using these identified keywords, we guide the LLM [8] to formulate an enhancement directive for the image content (e.g., the room's contentment and enjoyment). (iii) Finally, the LLM combines the original text description with the enhancement directive to generate the rich emotional prompt for the subsequent process. We present the details of the reasoning process in the Supp.

**Voting ensemble mechanism.** People may experience a variety of emotions when looking at the same image because they have different perceptual preferences, cognitive styles, and aesthetic tendencies. Consequently, rather than solely relying on VGG-based emotional distribution estimation, a voting ensemble mechanism is designed to simulate people viewing images from different perspectives, aggregating votes to accurately analyze the evoked emotions. Specifically, the emotional distribution is estimated separately using low-level features [3], [34] and mid-level features [35], [76]: (i) **Color.** Color histograms are independently computed for each of the RGB color channels in images. (ii) **Texture.** The gray-level co-occurrence matrix (GLCM) [34] is computed, followed by the extraction of handcrafted features. (iii) **Style.** A distribution estimator  $\varphi$  extracts style features (Sec. IV-B). (iv) **Patch.** The image is randomly cropped [76] and the patch features are extracted. Consequently,  $\phi^j(I^{\text{out}})_i$  is defined as the value of the  $i$ -th category within the distribution estimated by the  $j$ -th classifier:

$$\varphi(I^{\text{out}})_i = \frac{\sum_j \phi^j(I^{\text{out}})_i \times w^j}{\sum_j w^j}, \quad (21)$$

where ensemble weights  $w^j$  are the accuracies of classifier models from each perspective, evaluated on a hold-out validation set. Then, the emotional distribution loss  $\mathcal{L}_{\text{ed}}$  (Eq. 9) is adopted to learn high-dimensional emotional cues effectively.

### C. Emotional Reflection Strategy

As demonstrated in Fig. 3 (d), the sentiment metric loss and anchor-based sentiment loss are reused from AIF-B (Sec. IV-C). To effectively learn both emotional and artistic priors, sentiment metric loss  $\mathcal{L}_{\text{sm}}$  is combined with the style loss  $\mathcal{L}_s$  to improve the emotional reflection of the image decoder. Additionally, anchor-based sentiment loss  $\mathcal{L}_{\text{as}}$  is maintained to ensure the synthesized results evoke specific emotional responses.

**Emotion reflection decoder.** Observing that images synthesized by the pre-trained image decoder [6] cannot effectively produce colors and textures that align with the intended emotions, AIF-D enhances the emotional understanding of the decoder to appropriately present visually compelling artistic features. Specifically, sentiment vectors  $[V^{\text{sed}}, V^{\text{pos}}, V^{\text{rel}}, V^{\text{neg}}]$  are extracted from the images corresponding to text descriptions  $[T^{\text{sed}}, T^{\text{pos}}, T^{\text{rel}}, T^{\text{neg}}]$ . Defining the emotional distance as  $F_{\text{sw}}(V_i, V_j) = \frac{\|V_i - V_j\|^2}{F_{\text{dis}}(V_i, V_j)}$ , the sentiment loss  $\mathcal{L}_{\text{sm}}$  (Eq. 11)






Fig. 4. Visualization of applying the texture mapping loss across different blocks of the image decoder. (a) User-provided content images. (b) Text descriptions that reflect thoughts and feelings. (c)-(f) Results from applying the texture mapping loss at progressively later decoder blocks. (g) Results of AIF-D with appropriate artistic style and content consistency by combining losses across blocks.

is calculated to capture the relationships between emotions. To create appropriate colors and textures that reflect intended emotions, style loss  $\mathcal{L}_s$  (Eq. (14)) is adopted to minimize the style difference between synthesized results and anchor images. As a result, the image decoder is fine-tuned by anchoring visually abstract emotions to visually concrete images through a combined loss:

$$\mathcal{L}_{ft} = \mathcal{L}_{sm} + \lambda \mathcal{L}_s, \quad (22)$$

where hyper-parameter is set as  $\lambda = 0.01$ . Note that this process is completed before the training of the noise prediction network. The parameters of the image decoder are then fixed to preserve the learned emotional and artistic priors.

**Anchor-based sentiment loss.** Additionally, the anchor-based sentiment loss (Eq. (12)) is adopted to minimize the distance between sentiment vectors of the synthesized results and their corresponding anchor images. Within the diffusion-model backbone, this remains effective in ensuring the synthesized results evoke specific emotional responses from human observers.

#### D. Visual Aesthetics Balance

As shown in Fig. 3 (d), the aesthetic loss  $\mathcal{L}_{ae}$  is redesigned for the diffusion-based backbone to address the overly-stylized effects of AIF-B (Sec. IV-D). The redesigned loss emphasizes balancing artistic style and content consistency while visualizing emotions as appropriate colors and textures. AIF-D defines aesthetic loss as a combination of the diffusion model loss  $\mathcal{L}_{dm}$  and the texture mapping loss  $\mathcal{L}_{tm}$ , detailed as follows:

**Diffusion model loss.** Guided by the text tokens carrying emotional cues  $Z^{\text{emo}}$  and context of image tokens  $Z^{\text{img}}$ , the MSE loss is adopted to minimize the discrepancy between the learned image distribution and the target distribution for arbitrary noised latent code:

$$\mathcal{L}_{dm} = \mathbb{E}_{t, Z_0^{\text{img}}, \epsilon \sim \mathcal{N}(0,1)} [\|\epsilon_t - \epsilon_\theta(Z_t^{\text{img}}, t, Z^{\text{emo}}, Z_0^{\text{img}})\|^2]. \quad (23)$$

**Texture mapping Loss.** The texture mapping loss is presented to leverage the learned priors and the extracted decoding features. Denote the denoised latent observation as:

$$Z' = (Z_t^{\text{img}} - \sqrt{1 - \alpha_t} \epsilon_\theta(Z_t^{\text{img}}, t, Z^{\text{emo}}, Z_0^{\text{img}})) / \sqrt{\alpha_t}. \quad (24)$$

Anchor images are then fed into the image encoder and decoder to extract multi-scale decoding features as guidance:

$$\mathcal{L}_{tm}^i = F_{\text{dif}}(\mathcal{D}_i(Z^{\text{acr}}), \mathcal{D}_i(Z')), \quad (25)$$

where  $F_{\text{dif}}(a, b) = \|\mu(a) - \mu(b)\|_2 + \|\sigma(a) - \sigma(b)\|_2$  measures the difference between decoding features,  $\mathcal{D}_i$  means features extracted from the  $i$ -th decoding block, and  $Z^{\text{acr}} = \mathcal{E}(I^{\text{acr}})$  represents encoded anchor tokens. To demonstrate the influence of applying the texture mapping loss at different decoder blocks, the synthesized results are visualized in Fig. 4. As the loss is applied to progressively later blocks in the decoder, the artistic features become stronger, accompanied by increased content distortion. Therefore, the texture mapping loss is calculated across multiple blocks to balance artistic style and content consistency:

$$\mathcal{L}_{tm} = \sum_i \gamma^{i+1} \mathcal{L}_{tm}^i, \quad (26)$$

where  $\gamma = 0.3$  is a hyper-parameter.

Consequently, the aesthetic loss in AIF-D is defined as:

$$\mathcal{L}_{ae} = \lambda_{dm} \mathcal{L}_{dm} + \lambda_{tm} \mathcal{L}_{tm}, \quad (27)$$

where  $\lambda_{dm} = 1$  and  $\lambda_{tm} = 0.001$  denote the weights for losses.

#### E. Learning and training details

**Losses.** AIF-D is trained by combining the emotional distribution loss  $\mathcal{L}_{ed}$ , anchor-based sentiment loss  $\mathcal{L}_{as}$ , and aesthetic loss  $\mathcal{L}_{ae}$  into the full objective loss for training the noise prediction network:

$$\mathcal{L}_{\text{total}} = \lambda_{ed} \mathcal{L}_{ed} + \lambda_{as} \mathcal{L}_{as} + \lambda_{ae} \mathcal{L}_{ae} \quad (28)$$

where  $\lambda_{ed} = 10$ ,  $\lambda_{as} = 10$ , and  $\lambda_{ae} = 1$  are hyper-parameters, which are determined empirically.

**Training details.** The image decoder in AIF-D is first fine-tuned for only 2 epochs with a batch size of 4, a strategy similar to Imagic [7]. Subsequently, the parameters of the image decoder are fixed, and the noise prediction network is trained for 4 epochs with a batch size of 2, requiring approximately 68 hours with 2 NVIDIA GeForce 3090 GPUs. To supervise

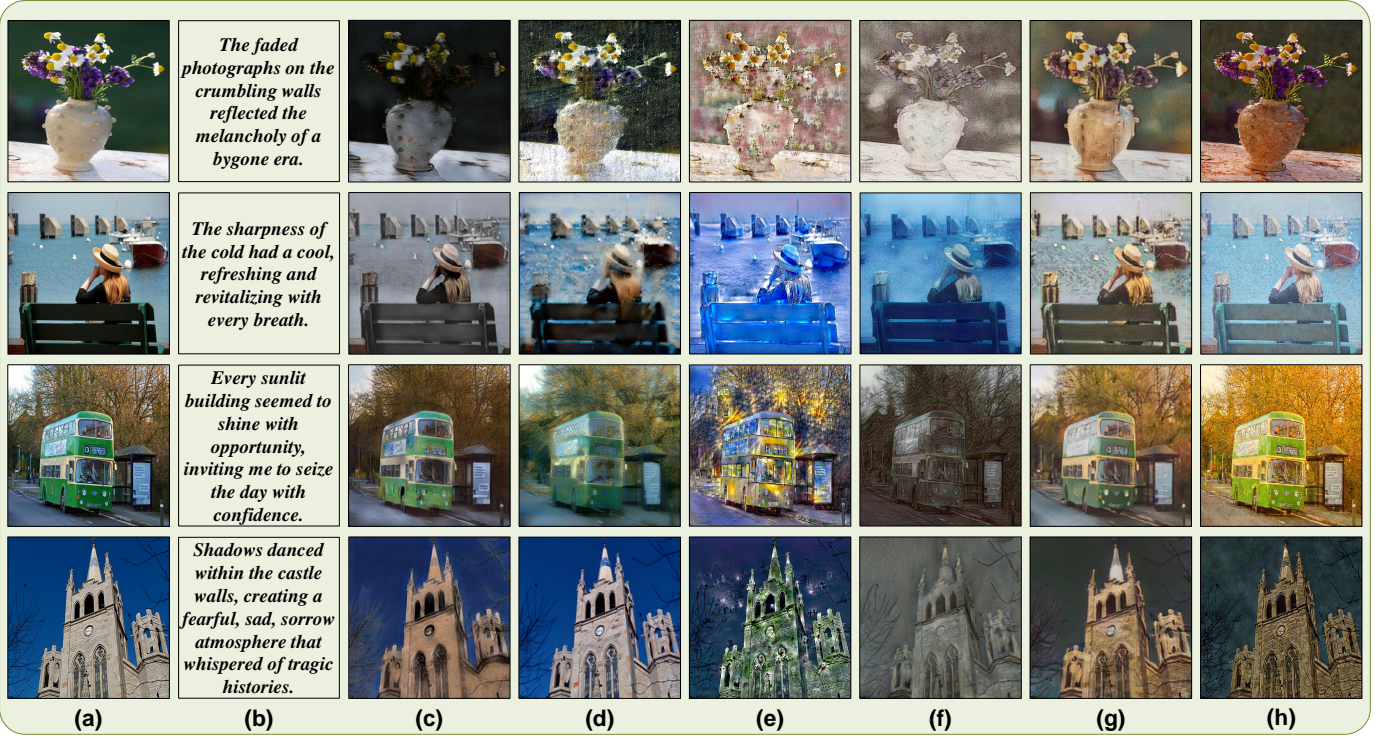


Fig. 5. Qualitative comparison results with state-of-the-art methods. (a) User-provided content images. (b) Text descriptions that reflect thoughts and feelings. (c) ManiGAN [16] (d) DiffusionCLIP [21]. (e) CLIPstyler [22]. (f) CLVA [13]. (g) AIF-B [2]. (h) AIF-D.

the synthesizing process conditioned on text descriptions, the classifier-free guidance [78] is applied. Adam optimizer [79] is employed with learning rates of  $1 \times 10^{-6}$  for the image decoder and the noise prediction network.

## VI. EXPERIMENTS

### A. Quantitative evaluation metrics

To conduct a thorough evaluation of the model’s performance in the AIF task, we utilize four quantitative metrics: (i) **Structural Similarity Index Measure (SSIM)**. Following CLVA [13], we use SSIM [80] to assess whether the synthesized images have similar visual content to the content images, evaluating whether the synthesized results preserve the content consistency. (ii) **Shallow Style Difference (SSD)**. Since the shallow layers of the VGG network extract color and texture features that are more aligned with AIF objectives, we calculate the style difference [53] between these layers to measure the statistical distribution of visual features between synthesized results and anchor images. (iii) **Sentiment Gap (SG)**. Following the preliminary work [2], we extract sentiment vectors of synthesized results and anchor images according to Yang *et al.* [69]. Their Euclidean distance is further calculated to determine whether synthesized images could evoke specific emotional responses. (iv) **Ensemble Accuracy (EAcc)**. We extend the Accuracy metric in our preliminary work [2], separately estimating emotional distributions from color, texture, style, and patch features. By ensembling these estimation results, we comprehensively evaluate whether the image filter reflects the intended emotions.

TABLE I  
QUANTITATIVE EXPERIMENT RESULTS.  $\uparrow$  ( $\downarrow$ ) MEANS HIGHER (LOWER) IS BETTER. THROUGHOUT THE PAPER, BEST PERFORMANCES ARE HIGHLIGHTED IN **BOLD**.

Methods	SSIM (%) $\uparrow$	SSD $\downarrow$	SG (%) $\downarrow$	EAcc (%) $\uparrow$
Comparison with state-of-the-art methods				
ManiGAN	50.72	1.4970	1.6589	33.58
DiffusionCLIP	53.05	1.9873	1.7095	32.71
CLIPstyler	52.49	1.9717	1.5676	34.19
CLVA	50.30	1.2232	1.6707	32.32
AIF-B (Ours)	56.15	1.1739	1.3881	36.02
AIF-D (Ours)	<b>57.74</b>	<b>1.1123</b>	<b>1.3416</b>	<b>37.48</b>
Ablation study of AIF-D				
W/o CPM	13.48	2.5374	1.8421	31.19
W/o IER	57.54	1.1251	1.3662	35.87
W/o VEM	57.32	1.1173	1.3949	34.18
W/o ERD	55.75	1.2356	1.4639	35.21
W/o TML	56.91	1.2824	1.4312	36.23

### B. Comparison with state-of-the-art methods

As the AIF task is a recently proposed task with limited closely related work, we compare our results with relevant text-driven image editing methods (*i.e.*, ManiGAN [16] and DiffusionCLIP [21]) and style transfer methods (*i.e.*, CLIPstyler [22] and CLVA [13]).

**Qualitative comparisons.** We show visual quality comparisons with the aforementioned methods in Fig. 5. Among these methods, ManiGAN [16] tends to distort content details (*e.g.*, the first row, the flowers in the vase become blurred); DiffusionCLIP [21] fails to preserve the overall structure of the






Fig. 6. Ablation study results with different variants of AIF-D. (a) User-provided content images. (b) Text descriptions that reflect thoughts and feelings. (c) W/o CPM. (d) W/o IER. (e) W/o VEM. (f) W/o ERD. (g) W/o TML. (h) AIF-D.

TABLE II

USER STUDY RESULTS. AIF-D OUTPERFORMS RELEVANT METHODS WITH THE HIGHEST SCORE.

Methods	EPS (%) $\uparrow$	EFS $\uparrow$	FES (%) $\uparrow$
ManiGAN	7.63	3.01	13.24
DiffusionCLIP	11.72	3.35	11.88
CLIPstyler	12.84	3.37	6.19
CLVA	9.78	3.03	9.30
AIF-B (Ours)	18.52	3.45	16.44
AIF-D (Ours)	<b>39.51</b>	<b>3.66</b>	<b>42.95</b>

content (e.g., the second row, the whole photo seems gritty); CLIPstyler [22] tends to over-stylize the content (the third row, the filter creates an unnaturally radiant golden glow); CLVA [13] generally synthesizes colorless images, making it difficult to accurately convey emotions (e.g., the fourth row, the castle appears shrouded in a layer of gray haze); AIF-B [2] can partially reflect emotions from text descriptions but struggles with effectively balancing artistic style and content consistency, resulting in blurry details (e.g., the fourth row, the texture on the castle walls looks irregular and distorted). In contrast, AIF-D demonstrates superior content consistency and emotional fidelity, advancing AIF models towards deeper emotional reflection.

**Quantitative comparisons.** We show comprehensive quantitative results in Tab. I where AIF-D achieves the highest scores across all four metrics. This demonstrates that AIF-D outperforms state-of-the-art methods. Specifically, AIF-D effectively preserves the overall structure of the content image (SSIM), exhibits appropriate filter-like visual effects (SSD),

TABLE III

HYPER-PARAMETER SENSITIVITY ANALYSIS FOR AIF-D MODEL.

Methods	SSIM (%) $\uparrow$	SSD $\downarrow$	SG (%) $\downarrow$	EAcc (%) $\uparrow$
AIF-D ( $\gamma = 0.1$ )	57.70	1.2057	1.4575	36.64
AIF-D ( $\gamma = 1.0$ )	57.03	1.1679	1.3694	37.04
AIF-D ( $\gamma = 3.0$ )	55.11	1.3142	1.6127	36.02
AIF-D ( $\gamma = 0.3$ )	<b>57.74</b>	<b>1.1123</b>	<b>1.3416</b>	<b>37.48</b>

evokes specific emotions of human observers (SG), and reflects emotions to images more accurately (EAcc).

**User study.** In addition to qualitative and quantitative comparisons, we further conduct three user studies to evaluate the effectiveness of AIF models. (i) **Emotional Preference Scores (EPS).** We conduct a user study to determine whether images synthesized by AIF model are more favored by human observers compared to those produced by relevant state-of-the-art methods. In this study, each participant is shown a combination of a content image, an emotional text, and six synthesized images. They are asked to select the synthesized image that they feel best matches the emotional text descriptions. We denote the rate of selection as the EPS for each method. (ii) **Emotional Fidelity Scores (EFS).** We investigate the emotional shift from the content images to the synthesized ones by asking participants to categorize the emotions of each synthesized image. The EFS ranging from 5 to 1 are calculated according to the emotional distance between the content images and synthesized images in Mikel’s wheel [63], where higher scores mean closer categories with respect to the labels of corresponding text descriptions. (iii) **Filter-like Effect**






Fig. 7. Application. **Left:** Retouching photography. By leveraging text descriptions that reflect users’ thoughts when taking the photo, AIF models can retouch these images to convey their intended emotions. **Right:** Social platforms. When sharing experiences on social platforms, users can select results that accurately reflect their emotions, thereby personalizing their content and attracting more followers.

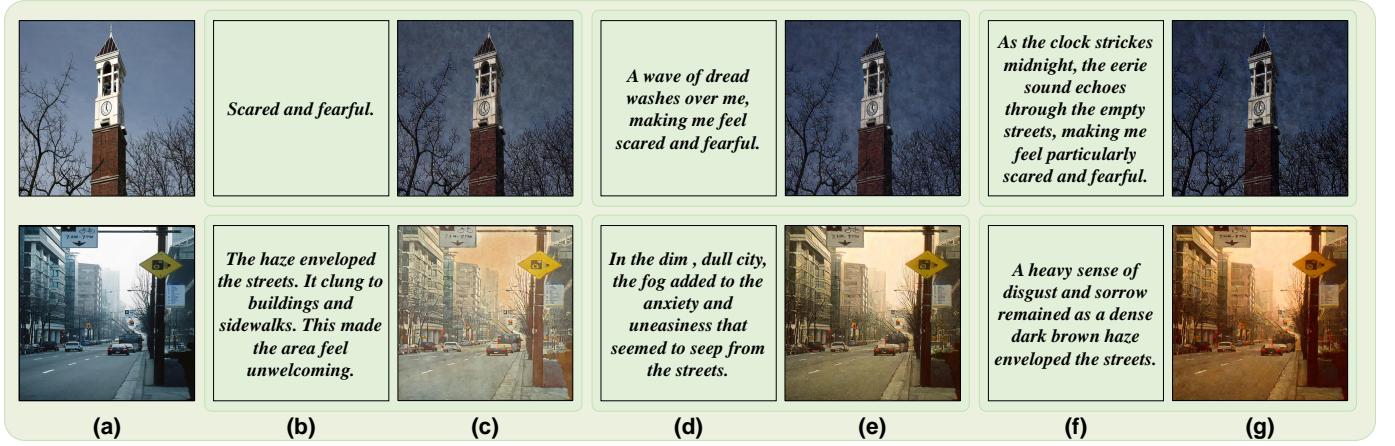


Fig. 8. Robustness testing. **Top:** Examples with varying levels of descriptive detail. **Bottom:** Examples with varying levels of emotional intensity. (a) User-provided content images. (b, d, f) Text descriptions that reflect thoughts and feelings. (c, e, g) Robust results produced by AIF-D.

**Scores (FES).** To further evaluate the visual effects of AIF models, we ask users to judge which image has the strongest filter-like effect. Participants are shown six synthesized images from comparison methods and are asked to select the one that appears most like the result of applying a filter. The FES for each method represents the selection rate. These experiments are conducted on Amazon Mechanical Turk (AMT), where 200 samples from the testing set of the AIF dataset are randomly selected, and experiment results are polled independently by 50 volunteers. We present the EPS, EFS, and FES results in Tab. III, where the highest scores indicate the subjective advantages of AIF-D. To further measure inter-rater agreement, we calculate Fleiss’s Kappa [81], resulting in 0.60 and 0.63 for EPS and FES, indicating a substantial level of consistency among raters.

### C. Ablation study

We discard various modules and create five variants to study the impacts of AIF-D’s modules and designed losses. The evaluation scores and synthesized images of the ablation study are shown in Tab. I and Fig. 6, respectively.

**W/o CPM (Content Preservation Module).** In this ablation, we remove the content preservation module. As a result,

this ablation cannot preserve the overall structure without corresponding content images as cues (lower SSIM score). As shown in the first row of Fig. 6, synthesized images are completely irrelevant to the user-provided content images and are therefore unqualified.

**W/o IER (In-depth Emotional Reasoning).** We discard the in-depth emotional reasoning module. Due to the absence of emotional and filter effect prompts, AIF-D cannot leverage the emotional cues from the LLM to infer deeper emotions (higher SG score and lower EAcc score). As shown in the second row of Fig. 6, the overall color tone of the image appears dull, failing to convey the emotion of fear.

**W/o VEM (Voting Ensemble Mechanism).** We remove the voting ensemble mechanism. Consequently, AIF-D has difficulty in comprehensively understanding ambiguous and subjective emotions (reduced EAcc score). As shown in the third row of Fig. 6, the overall color tone of the image is on the darker side, failing to convey the disgusting feeling mentioned in the text description effectively.

**W/o ERD (Emotional Reflection Decoder).** We replace the emotional reflection image decoder with the vanilla image decoder. Therefore, this variant struggles to anchor emotions

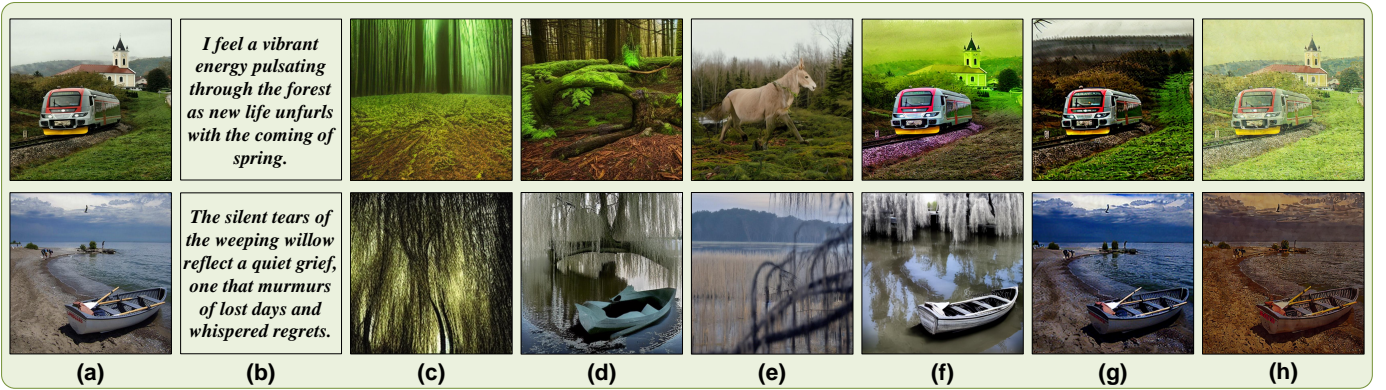


Fig. 9. Qualitative comparison results with general image editing models based on the diffusion approach. (a) User-provided content images. (b) Texts that reflect thoughts and feelings. (c) Stable Diffusion [6]. (d) ControlNet [27]. (e) SDEdit [82]. (f) InstructPix2Pix [83]. (g) Imagic [77]. (h) AIF-D.

to colors and textures with appropriate artistic representations (higher SG score). As shown in the fourth row of Fig. 6, there are numerous white and black artifacts scattered across the image, damaging the image emotional representation.

**W/o TML (Texture Mapping Loss).** During the training of the noise prediction network, we utilize the standard style loss [53] instead of the proposed texture mapping loss. As a result, the model excessively applies artistic modifications to user-provided content images (higher SSD score). As shown in the second row of Fig. 6, the room presents exaggerated contrast, leading to overly-stylized effects.

#### D. Hyper-parameter sensitivity analysis.

We further investigate the influence of the crucial hyper-parameter  $\gamma$  (introduced in Eq. (26)), which controls the balance between artistic style and content consistency. As Tab. III demonstrates that  $\gamma = 0.3$  achieves the best quantitative results, we adopt this value for the AIF-D model.

### VII. DISCUSSION

#### A. Application

The AIF models are typically applied in two scenarios.

**Retouching photography.** With daily photography becoming a ubiquitous practice, there is a growing expectation for photos to have artistic qualities, be visually pleasing, and reflect the photographer’s mood. Given users’ written thoughts, AIF models could retouch their photos to be visually compelling and reflect what they are thinking. We show this application in Fig. 7 (left), where the user requests the model to enhance their photograph by applying affective filters that reflect the emotion they experienced at the moment of capture.

**Social platforms.** People share posts with texts and corresponding images to express their feelings on social platforms. AIF models could personalize their content, potentially enabling influencers to stand out from the crowd and evoke specific emotional responses from their followers. As demonstrated in Fig. 7 (right), several affective images are provided as options, recognizing that emotions are inherently subjective and ambiguous. This allows users to select the result that most accurately represents their intended visually-abstract emotions.

#### B. Robustness testing of AIF-D

To evaluate the robustness of AIF-D’s text understanding with diverse user input (e.g., varying levels of details and emotional intensity), we conduct the following experiments.

**Levels of detail.** As shown in Fig. 8 (top), we present examples with consistent emotional intensity but varying levels of detail. As a result, AIF-D effectively conveys emotions (particularly scared and fearful), regardless of the levels of descriptive detail.

**Levels of intensity.** As shown in Fig. 8 (bottom), we maintain descriptive detail but change emotional intensity. Consequently, AIF-D successfully produces results that progressively heighten the sense of emotions (disgust and discomfort).

#### C. Differences from general image editing

Reflecting visually-abstract emotions from text descriptions to visually concrete images presents a significant challenge. Existing general image editing techniques achieve impressive results in their respective domains but struggle to meet the objectives of the AIF task. Specifically, Stable Diffusion [6] and SDEdit [82] struggle to preserve the overall structure of content images; ControlNet [27], InstructPix2Pix [83], and Imagic [77] are unable to understand emotional cues to synthesize appropriate artistic representation. We illustrate their unsuitability in Fig. 9.

### VIII. CONCLUSION

In this paper, we introduce the task of Affective Image Filter (AIF), which requires the model to understand visually-abstract emotions from text descriptions and reflect them into visually-concrete images with appropriate colors and textures. We first introduce the AIF dataset and formulate the AIF model. Next, we present AIF-B, built upon a multi-modal transformer architecture. As an initial attempt, AIF-B leverages external emotional cues, proposing the emotional reflection loss, and designing the visual aesthetic loss. To advance AIF models towards deeper emotional reflection, we propose AIF-D, a model built upon a pre-trained diffusion model with a content preservation module for preserving high-frequency detail. By ensembling emotional cues, improving emotional reflection strategy, and redesigning visual aesthetic

balance, AIF-D achieves state-of-the-art performance across four quantitative and three user study experiments. Finally, we discuss applications, robustness, and differences from general image editing techniques to highlight the value and potential of AIF models.

#### ACKNOWLEDGEMENT

This work is supported by National Natural Science Foundation of China (Grant No. 62136001) and Beijing Municipal Science & Technology Commission, Administrative Commission of Zhongguancun Science Park (Grant No. Z241100003524012). We thank vivo for providing the GPU resources and thank Qingnan Fan and Jinwei Chen for their insightful discussions and valuable feedback.

#### REFERENCES

- [1] E. Kross, P. Verduyn, G. Sheppes, C. K. Costello, J. Jonides, and O. Ybarra, "Social media and well-being: Pitfalls, progress, and next steps," *Trends in Cognitive Sciences*, vol. 25, no. 1, pp. 55–66, 2021.
- [2] S. Weng, P. Zhang, Z. Chang, X. Wang, S. Li, and B. Shi, "Affective image filter: Reflecting emotions from text to images," in *ICCV*, 2023.
- [3] J. Machajdik and A. Hanbury, "Affective image classification using features inspired by psychology and art theory," in *ACM MM*, 2010.
- [4] S. Zhao, H. Yao, Y. Yang, and Y. Zhang, "Affective image retrieval via multi-graph learning," in *ACM MM*, 2014.
- [5] S. Mohammad, "Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 english words," in *ACL*, 2018.
- [6] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *CVPR*, 2022.
- [7] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [8] Z. Yang, L. Li, K. Lin, J. Wang, C.-C. Lin, Z. Liu, and L. Wang, "The dawn of LMMs: Preliminary explorations with GPT-4V (ision)," *arXiv preprint arXiv:2309.17421*, 2023.
- [9] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," in *NeurIPS*, 2022.
- [10] Z. Li, G. Chen, R. Shao, D. Jiang, and L. Nie, "Enhancing the emotional generation capability of large language models via emotional chain-of-thought," *arXiv preprint arXiv:2401.06836*, 2024.
- [11] S. Weng, H. Wu, Z. C. Chang, J. Tang, S. Li, and B. Shi, "L-CoDe: Language-based colorization using color-object decoupled conditions," in *AAAI*, 2022.
- [12] Z. Chang, S. Weng, P. Zhang, Y. Li, S. Li, and B. Shi, "L-CAD: Language-based colorization with any-level descriptions using diffusion priors," in *NeurIPS*, 2023.
- [13] T.-J. Fu, X. E. Wang, and W. Y. Wang, "Language-driven artistic style transfer," in *ECCV*, 2022.
- [14] G. Couairon, A. Grechka, J. Verbeek, H. Schwenk, and M. Cord, "Flexit: Towards flexible semantic image translation," in *CVPR*, 2022.
- [15] N. Tumanyan, M. Geyer, S. Bagon, and T. Dekel, "Plug-and-play diffusion features for text-driven image-to-image translation," in *CVPR*, 2023.
- [16] B. Li, X. Qi, T. Lukasiewicz, and P. H. Torr, "ManiGAN: Text-guided image manipulation," in *CVPR*, 2020.
- [17] S. Weng, W. Li, D. Li, H. Jin, and B. Shi, "Conditional image repainting via semantic bridge and piecewise value function," in *ECCV*, 2020.
- [18] L. Zhang, Q. Chen, B. Hu, and S. Jiang, "Text-guided neural image inpainting," in *ACM MM*, 2020.
- [19] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *ICML*, 2021.
- [20] A. Q. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, "Glide: Towards photorealistic image generation and editing with text-guided diffusion models," in *ICML*, 2022.
- [21] G. Kim, T. Kwon, and J. C. Ye, "DiffusionCLIP: Text-guided diffusion models for robust image manipulation," in *CVPR*, 2022.
- [22] G. Kwon and J. C. Ye, "CLIPstyler: Image style transfer with a single text condition," in *CVPR*, 2022.
- [23] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *ICLR*, 2021.
- [24] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *NeurIPS*, 2020.
- [25] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," *arXiv preprint arXiv:2204.06125*, 2022.
- [26] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans *et al.*, "Photorealistic text-to-image diffusion models with deep language understanding," in *NeurIPS*, 2022.
- [27] L. Zhang, A. Rao, and M. Agrawala, "Adding conditional control to text-to-image diffusion models," in *ICCV*, 2023.
- [28] C. Mou, X. Wang, L. Xie, J. Zhang, Z. Qi, Y. Shan, and X. Qie, "T2I-Adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models," *arXiv preprint arXiv:2302.08453*, 2023.
- [29] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, "Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation," in *CVPR*, 2023.
- [30] N. Kumari, B. Zhang, R. Zhang, E. Shechtman, and J.-Y. Zhu, "Multi-concept customization of text-to-image diffusion," in *CVPR*, 2023.
- [31] A. Hertz, R. Mokady, J. Tenenbaum, K. Aberman, Y. Pritch, and D. Cohen-Or, "Prompt-to-prompt image editing with cross attention control," *arXiv*, 2022.
- [32] G. Parmar, K. K. Singh, R. Zhang, Y. Li, J. Lu, and J.-Y. Zhu, "Zero-shot image-to-image translation," *arXiv preprint arXiv:2302.03027*, 2023.
- [33] S. Zhao, X. Yao, J. Yang, G. Jia, G. Ding, T.-S. Chua, B. W. Schuller, and K. Keutzer, "Affective image content analysis: Two decades review and new perspectives," *TPAMI*, 2021.
- [34] S. Zhao, Y. Gao, X. Jiang, H. Yao, T.-S. Chua, and X. Sun, "Exploring principles-of-art features for image emotion recognition," in *ACM MM*, 2014.
- [35] J. Yang, D. She, Y.-K. Lai, P. L. Rosin, and M.-H. Yang, "Weakly supervised coupled networks for visual sentiment analysis," in *CVPR*, 2018.
- [36] X. Yao, D. She, H. Zhang, J. Yang, M.-M. Cheng, and L. Wang, "Adaptive deep metric learning for affective image retrieval and classification," *TMM*, 2020.
- [37] T. Rao, X. Li, and M. Xu, "Learning multi-level deep representations for image emotion classification," *NPL*, 2020.
- [38] J. Yang, D. She, M. Sun, M.-M. Cheng, P. L. Rosin, and L. Wang, "Visual sentiment prediction based on automatic discovery of affective regions," *TMM*, 2018.
- [39] L. Wang, G. Jia, N. Jiang, H. Wu, and J. Yang, "Ease: Robust facial expression recognition via emotion ambiguity-sensitive cooperative networks," in *ACM MM*, 2022.
- [40] D. Yang, Z. Chen, Y. Wang, S. Wang, M. Li, S. Liu, X. Zhao, S. Huang, Z. Dong, P. Zhai *et al.*, "Context de-confounded emotion recognition," in *CVPR*, 2023.
- [41] J. Wang, B. Xu, and Y. Zu, "Deep learning for aspect-based sentiment analysis," in *MLISE*, 2021.
- [42] X. Wang, J. Jia, J. Yin, and L. Cai, "Interpretable aesthetic features for affective image classification," in *ICIP*, 2013.
- [43] T. Rao, M. Xu, H. Liu, J. Wang, and I. Burnett, "Multi-scale blocks based image emotion classification using multiple instance learning," in *ICIP*, 2016.
- [44] D. Yang, S. Huang, S. Wang, Y. Liu, P. Zhai, L. Su, M. Li, and L. Zhang, "Emotion recognition for multiple context awareness," in *ECCV*, 2022.
- [45] D. Yang, K. Yang, H. Kuang, Z. Chen, Y. Wang, and L. Zhang, "Towards context-aware emotion recognition debiasing from a causal demystification perspective via de-confounded training," *PAMI*, 2024.
- [46] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *CVPR*, 2016.
- [47] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *ECCV*, 2016.
- [48] C. Li and M. Wand, "Precomputed real-time texture synthesis with markovian generative adversarial networks," in *ECCV*, 2016.
- [49] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *ICCV*, 2017.
- [50] Y. Li, C. Fang, J. Yang, Z. Wang, X. Lu, and M.-H. Yang, "Universal style transfer via feature transforms," in *NeurIPS*, 2017.
- [51] D. Chen, L. Yuan, J. Liao, N. Yu, and G. Hua, "StyleBank: An explicit representation for neural image style transfer," in *CVPR*, 2017.
- [52] J. Svoboda, A. Anoosheh, C. Osendorfer, and J. Masci, "Two-stage peer-regularized feature recombination for arbitrary image style transfer," in *CVPR*, 2020.



- [53] Y. Deng, F. Tang, W. Dong, C. Ma, X. Pan, L. Wang, and C. Xu, "StyTr2: Image style transfer with transformers," in *CVPR*, 2022.
- [54] S. Liu, T. Lin, D. He, F. Li, M. Wang, X. Li, Z. Sun, Q. Li, and E. Ding, "Adaattn: Revisit attention mechanism in arbitrary neural style transfer," in *ICCV*, 2021.
- [55] D. Y. Park and K. H. Lee, "Arbitrary style transfer with style-attentional networks," in *CVPR*, 2019.
- [56] X. Wu, Z. Hu, L. Sheng, and D. Xu, "Styleformer: Real-time arbitrary style transfer via parametric style composition," in *ICCV*, 2021.
- [57] D.-Y. Chen, "Artfusion: Arbitrary style transfer using dual conditional latent diffusion models," *arXiv preprint arXiv:2306.09330*, 2023.
- [58] S. Li, J. van de Weijer, T. Hu, F. S. Khan, Q. Hou, Y. Wang, and J. Yang, "StyLED: Prompt-embedding inversion for text-based editing," in *ICCV*, 2023.
- [59] Y. Zhang, N. Huang, F. Tang, H. Huang, C. Ma, W. Dong, and C. Xu, "Inversion-based style transfer with diffusion models," in *CVPR*, 2023.
- [60] S. Yang, H. Hwang, and J. C. Ye, "Zero-shot contrastive loss for text-guided diffusion image style transfer," in *ICCV*, 2023.
- [61] P. Achlioptas, M. Ovsjanikov, K. Haydarov, M. Elhoseiny, and L. J. Guibas, "ArtEmis: Affective language for visual art," in *CVPR*, 2021.
- [62] Y. Mohamed, F. F. Khan, K. Haydarov, and M. Elhoseiny, "It is okay to not be okay: Overcoming emotional bias in affective image captioning by contrastive data collection," in *CVPR*, 2022.
- [63] J. A. Mikels, B. L. Fredrickson, G. R. Larkin, C. M. Lindberg, S. J. Maglio, and P. A. Reuter-Lorenz, "Emotional category data on images from the international affective picture system," *BRM*, 2005.
- [64] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and psychological measurement*, 1960.
- [65] H. Caesar, J. R. R. Uijlings, and V. Ferrari, "Coco-stuff: Thing and stuff classes in context," in *CVPR*, 2018.
- [66] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *NAACL*, 2019.
- [67] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021.
- [68] S. Kullback and R. A. Leibler, "On information and sufficiency," *AoMS*, 1951.
- [69] J. Yang, D. She, Y.-K. Lai, and M.-H. Yang, "Retrieving and classifying affective images via deep metric learning," in *AAAI*, 2018.
- [70] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. N. Metaxas, "Stackgan++: Realistic image synthesis with stacked generative adversarial networks," *TPAMI*, 2018.
- [71] D. Y. Park and K. H. Lee, "Arbitrary style transfer with style-attentional networks," in *CVPR*, 2019.
- [72] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [73] R. Xiong, Y. Yang, D. He, K. Zheng, S. Zheng, C. Xing, H. Zhang, Y. Lan, L. Wang, and T. Liu, "On layer normalization in the transformer architecture," in *ICML*, 2020.
- [74] D. P. Kingma, "Auto-encoding variational bayes," in *ICLR*, 2014.
- [75] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017.
- [76] T. Park, J.-Y. Zhu, O. Wang, J. Lu, E. Shechtman, A. Efros, and R. Zhang, "Swapping autoencoder for deep image manipulation," in *NeurIPS*, 2020.
- [77] B. Kavar, S. Zada, O. Lang, O. Tov, H. Chang, T. Dekel, I. Mosseri, and M. Irani, "Imagic: Text-based real image editing with diffusion models," in *CVPR*, 2023.
- [78] J. Ho and T. Salimans, "Classifier-free diffusion guidance," in *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- [79] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [80] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *TIP*, 2004.
- [81] J. L. Fleiss, "Measuring nominal scale agreement among many raters," *Psychological bulletin*, 1971.
- [82] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon, "Sdedit: Guided image synthesis and editing with stochastic differential equations," *arXiv preprint arXiv:2108.01073*, 2021.
- [83] T. Brooks, A. Holynski, and A. A. Efros, "Instructpix2pix: Learning to follow image editing instructions," in *CVPR*, 2023.



**Peixuan Zhang** is a PhD student at the School of Artificial Intelligence, Beijing University of Posts and Telecommunications. His current research interests include affective computing and cross-modality (mainly language-based) content creation.



**Shuchen Weng** is a research scientist at Beijing Academy of Artificial Intelligence. He received the PhD degree from Peking University in 2024. His research interests include cross-modality (mainly language-based) content creation and manipulation.



**Jiajun Tang** is a research engineer at vivo. He received the PhD degree from Peking University in 2025. His research interests include cross-modal content creation and manipulation (mainly lighting-related), as well as 3D scene reconstruction.



**Si Li** received the Ph.D. degree from the Beijing University of Posts and Telecommunications in 2012. She is currently an Associate Professor with the School of Artificial Intelligence, Beijing University of Posts and Telecommunications. Her current research interests include multimodal artificial intelligence and machine learning.



**Boxin Shi** (Senior Member, IEEE) received the BE degree from the Beijing University of Posts and Telecommunications, the ME degree from Peking University, and the PhD degree from the University of Tokyo, in 2007, 2010, and 2013. He is currently a Boya Young Fellow Associate Professor (with tenure) and Research Professor at Peking University, where he leads the Camera Intelligence Lab. Before joining PKU, he did research with MIT Media Lab, Singapore University of Technology and Design, Nanyang Technological University, National Institute of Advanced Industrial Science and Technology, from 2013 to 2017. His papers were awarded as Best Paper, Runners-Up at CVPR 2024, ICCP 2015 and selected as Best Paper candidate at ICCV 2015. He is an associate editor of TPAMI/IJCV and an area chair of CVPR/ICCV/ECCV. He is a Distinguished Lecturer of APSIPA, a Distinguished Member of CCF, and a Senior Member of IEEE.

# Supplementary Materials of Towards Deeper Emotional Reflection: Crafting Affective Image Filters with Generative Priors

Peixuan Zhang<sup>†</sup>, Shuchen Weng<sup>†</sup>, Jiajun Tang, Si Li\* and Boxin Shi, *Senior Member, IEEE*

## A. In-depth emotional reasoning

As illustrated in Sec. V-B of the main paper, we leverage an LLM [1] and the CoT prompting to facilitate in-depth emotional reasoning. For better reproducibility of our results, an example of this reasoning process is illustrated in Fig. S1.



Fig. S1. Detailed illustration of in-depth emotional reasoning.

## B. Voting ensemble mechanism

As illustrated in Eq. (21) (Sec. V-B) of the main paper, a voting ensemble mechanism is developed that simulates people looking at images from different points of view and ensembles votes for accurately analyzing major evoked emotions. Specifically, distinct features are first extracted from synthesized

results, *e.g.*, color, texture, style, and patch features. These features are then fed into the corresponding classifier to estimate the emotional distribution independently. The backbone of these classifiers is a three-layer convolutional neural network. Each classifier is trained with the Kullback-Leibler (KL) loss function [2], enabling them to effectively analyze the emotional content of the images. The accuracy of each classifier is presented in Tab. S1. During training, the parameters of these classifiers are frozen, and their estimated distributions are ensembled to effectively learn high-dimensional emotional cues. The architecture of the voting ensemble mechanism is shown in Fig. S2.

TABLE S1  
THE ACCURACY (%) OF FOUR INDEPENDENT CLASSIFIERS (COLOR, TEXTURE, STYLE, AND PATCH) AND THE COMBINED ENSEMBLE RESULT.

Method	Color	Texture	Style	Patch	Ensemble
ManiGAN	30.99	27.21	27.77	30.47	33.58
DiffusionCLIP	30.39	28.62	24.59	30.75	32.71
CLIPstyler	30.99	29.61	26.40	30.39	34.19
CLVA	30.17	30.05	25.64	30.59	32.32
AIF-B (Ours)	31.12	31.72	29.96	31.77	36.02
AIF-D (Ours)	<b>32.92</b>	<b>32.12</b>	<b>30.87</b>	<b>32.57</b>	<b>37.48</b>

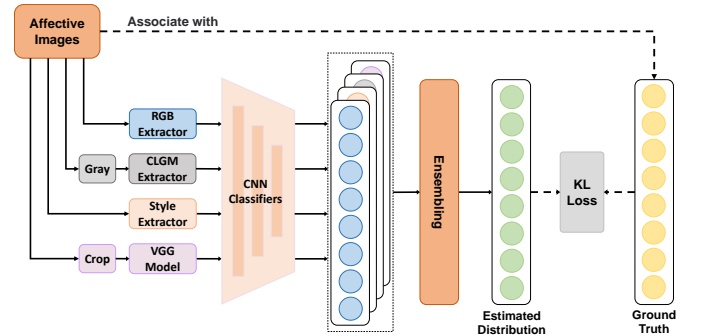


Fig. S2. The architecture of the voting ensemble mechanism.

## C. Rater instruction details

As mentioned in Sec. VI-B of the main paper, we conduct three user studies to evaluate the effectiveness of AIF models. To illustrate the detailed instructions, we present the user study interface in Fig. S3.

P. Zhang and S. Li are with the School of Artificial Intelligence, Beijing University of Posts and Telecommunications, Beijing 100876, China. Email: {pxzhang, lisi}@bupt.edu.cn.

S. Weng is with the Beijing Academy of Artificial Intelligence, Beijing 100083, China. Email: {scweng}@baai.ac.cn.

J. Tang and B. Shi are with the State Key Laboratory of Multimedia Information Processing and National Engineering Research Center of Visual Technology, School of Computer Science, Peking University, Beijing 100871, China. Email: {jiajun.tang, shiboxin}@pku.edu.cn.

<sup>†</sup> Equal contribution. \* Corresponding author.

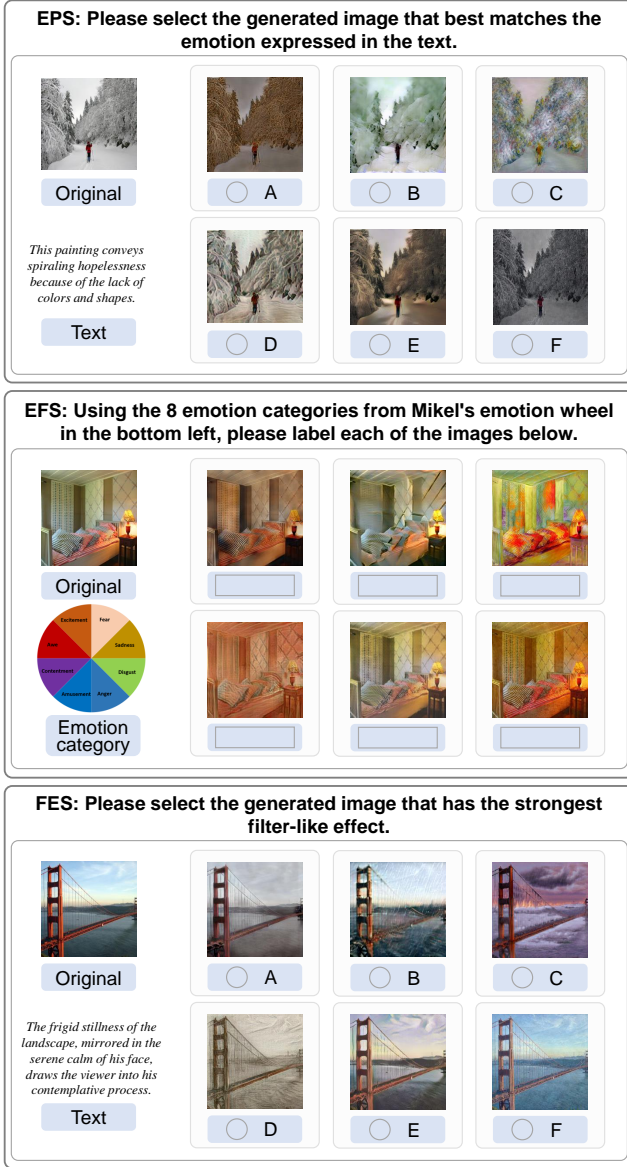


Fig. S3. The user interface for user study experiments.

#### D. Failure cases

We present two failure cases that the AIF models encounter. **Applying counterintuitive content.** The visual content of an image plays a critical role in evoking specific emotional responses from human observers. AIF models still face challenges in producing convincing results when the visual content of the provided image conflicts with the intended emotion of text descriptions. As shown in Fig. S4 (top), AIF-D fails to accurately reflect the emotion of “sorrow” in the image where a girl is smiling and struggles to reflect “peacefulness” in the image depicting a burning grassland.

**Reflecting contrastive emotions.** AIF models also face challenges in reflecting contrasting emotions within varied and complex text descriptions, leading to difficulties in creating appropriate affective image filters. As illustrated in Fig. S4 (bottom), text descriptions that combine opposing emotions (e.g., “triumph/defeat” or “awe/tension”) pose significant



Fig. S4. Failure cases. **Top:** Examples where the content image conflicts with intended emotions. **Bottom:** Examples where the text description includes opposing emotions.

challenges for AIF-D, resulting in inaccuracies in reflecting the intended emotions of users.

#### E. Additional qualitative results

We show additional comparison results of AIF models with relevant text-driven image editing methods (i.e., ManiGAN [3] and DiffusionCLIP [4]) and style transfer methods (i.e., CLIPstyler [5] and CLVA [6]). The qualitative results are presented in Fig. S5. More qualitative ablation study results are shown in Fig. S6, whose details are presented in Sec. VI of the main paper. More application results are shown in Fig. S7. With different text descriptions, users can create visually compelling photos that evoke a variety of emotional responses.

#### F. Dataset samples

As illustrated in Sec. III-A of the main paper, a large-scale dataset tailored for the AIF task is collected and processed, including abundant aesthetic images and corresponding text descriptions associated with the closest emotional category in Mikel’s wheel [7]. More samples of the AIF dataset are shown in Fig. S8 to inspire relevant research.

#### REFERENCES

- [1] Z. Yang, L. Li, K. Lin, J. Wang, C.-C. Lin, Z. Liu, and L. Wang, “The dawn of LMMs: Preliminary explorations with GPT-4V (ision),” *arXiv preprint arXiv:2309.17421*, 2023.
- [2] S. Kullback and R. A. Leibler, “On information and sufficiency,” *AoMS*, 1951.
- [3] B. Li, X. Qi, T. Lukasiewicz, and P. H. Torr, “ManiGAN: Text-guided image manipulation,” in *CVPR*, 2020.
- [4] G. Kim, T. Kwon, and J. C. Ye, “DiffusionCLIP: Text-guided diffusion models for robust image manipulation,” in *CVPR*, 2022.
- [5] G. Kwon and J. C. Ye, “CLIPstyler: Image style transfer with a single text condition,” in *CVPR*, 2022.
- [6] T.-J. Fu, X. E. Wang, and W. Y. Wang, “Language-driven artistic style transfer,” in *ECCV*, 2022.
- [7] J. A. Mikels, B. L. Fredrickson, G. R. Larkin, C. M. Lindberg, S. J. Maglio, and P. A. Reuter-Lorenz, “Emotional category data on images from the international affective picture system,” *BRM*, 2005.
- [8] S. Weng, P. Zhang, Z. Chang, X. Wang, S. Li, and B. Shi, “Affective image filter: Reflecting emotions from text to images,” in *ICCV*, 2023.





Fig. S5. Additional qualitative comparison results with state-of-the-art methods. (a) User-provided content images. (b) Texts that reflect thoughts and feelings. (c) ManiGAN [3]. (d) DiffusionCLIP [4]. (e) CLIPstyler [5]. (f) CLVA [6]. (g) AIF-B [8]. (h) AIF-D.



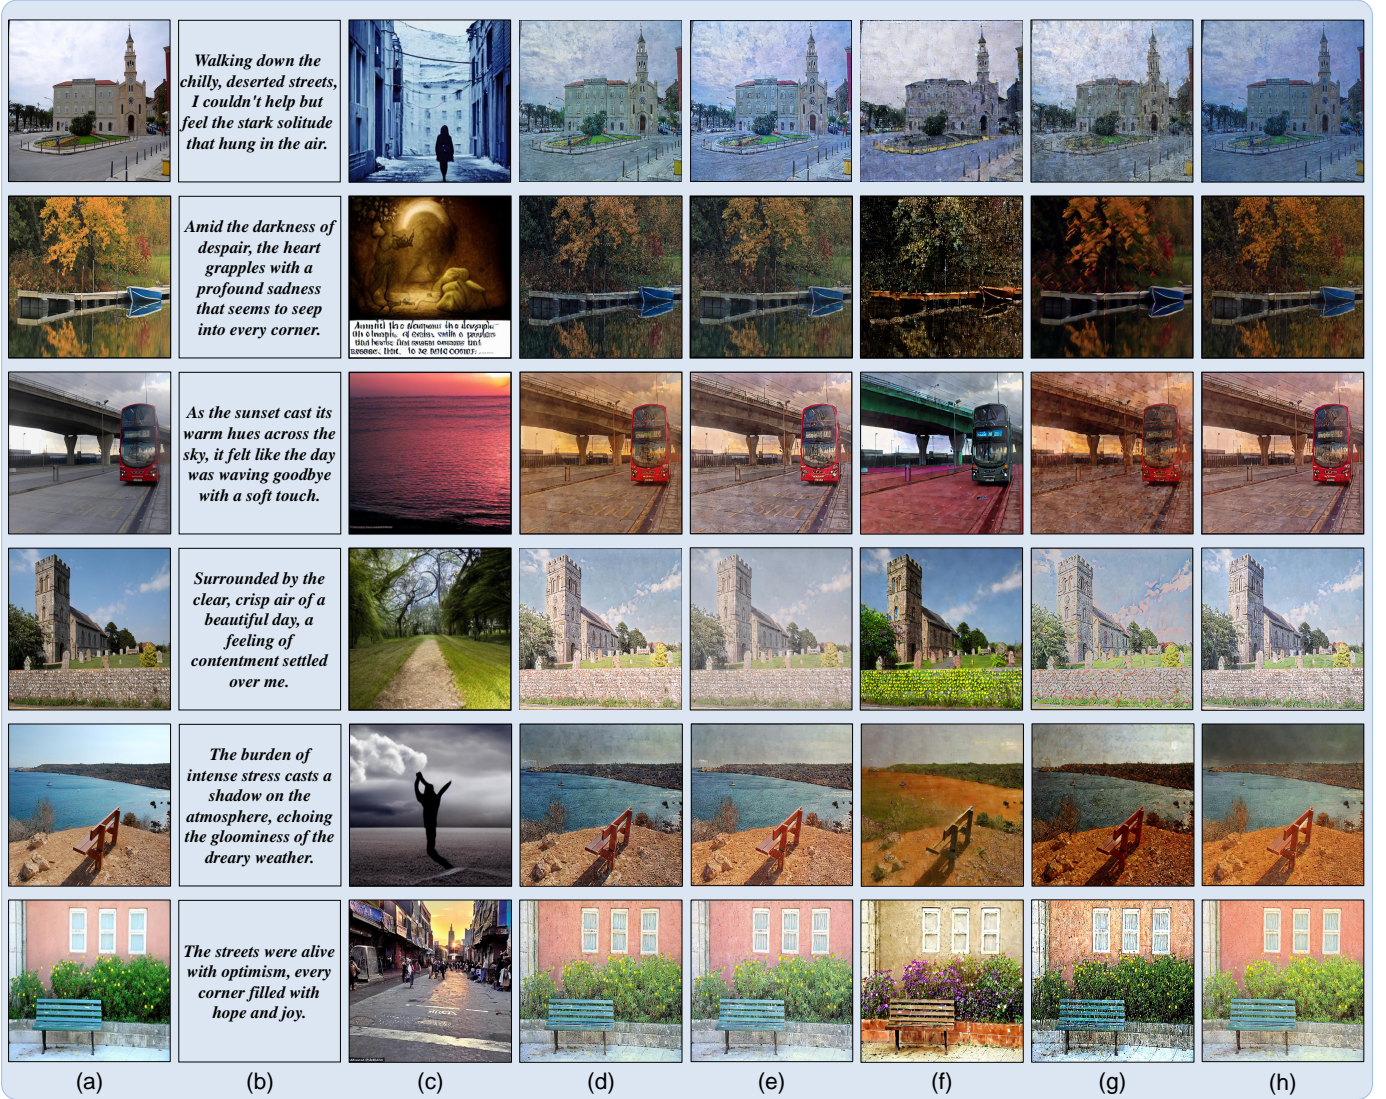


Fig. S6. Additional ablation study results with different variants of AIF-D. (a) User-provided content images. (b) Texts that reflect thoughts and feelings. (c) W/o CPM. (d) W/o IER. (e) W/o VEM. (f) W/o ERD. (g) W/o TML. (h) AIF-D.



Fig. S7. Additional applications. With different written thoughts, users could produce visually compelling photos that evoke a range of emotional responses.



<i>The bight whites and shades of the sand blend well with the light blue of the water.</i>	<i>This horrid boring and very sad and gloomy picture of the ocean and beach is sullen.</i>	<i>This sad cloudy skies and barren beach looks awkward and down beat.</i>	<i>The clouds are ruining the nice view.</i>	<i>Nobody on the beach and being able to appreciate and take in the feeling of amazement is inspired.</i>	<i>This painting looks very dreary and extremely upsetting.</i>	
Excitement	Sadness	Sadness	Sadness	Awe	Fear	
<i>I love the composition and how dark this is. The light reflections in the air is fantastic.</i>	<i>The way the light above the building glows in the dark is eerie.</i>	<i>the long dark road with the giant mansion in the dead of night is the start to horror movies.</i>	<i>Spooky nighttime sky, dark bushes, ominous mist and shadows created by the lights.</i>	<i>There is a green cloud in the background, almost like a coming darkness.</i>	<i>There's a lone carriage despite how large the building is, creating a feeling of loneliness and isolation.</i>	
Excitement	Fear	Fear	Fear	Fear	Fear	
<i>Seeing the smoke in the air makes me think of all the future air pollution.</i>	<i>The boat looks serene and fun out on the water.</i>	<i>The colors make this feel warm and pleasant and the water appears calm and peaceful.</i>	<i>As the viewer approaches the port, a sense of disgust can be felt by the smoke and dirty bay.</i>	<i>the smoke stacks show the industrialization of society ruining nature.</i>	<i>What looks like smoke in the distance, gives of an impression that there is a riot currently happening.</i>	
Anger	Awe	Contentment	Disgust	Anger	Fear	
<i>The tabby cat looks pleased with himself after causing some shenanigans earlier in the day.</i>	<i>The cat looks like he's smiling and it feels to me like he's in a playful mood.</i>	<i>I am in awe of this beautifully painted cat, especially the fur and the facial expression.</i>	<i>The cat seems sleepy, and the amount of detail in the fur makes it appear soft and well cared for.</i>	<i>I feel content as the painting is of a cat that seems to have a subtle smile across its face.</i>	<i>While the cat is sweet and beautiful, there is a bit of an aggression in his stance that is a bit scary.</i>	
Amusement	Amusement	Contentment	Contentment	Contentment	Fear	
<i>These sad flowers with yellows and reds in a glass vase look despondent.</i>	<i>The vase of flowers are beautifully painted and it makes me feel happy.</i>	<i>The flowers in the reflective vase over the stand appear majestic and elegant in a comforting manner.</i>	<i>The painting is amazing and the colors are pleasing to the eyes.</i>	<i>The drab colors of the wall just make everything in the room seem depressing.</i>	<i>The tone of the colors, especially on the walls behind the vase are quite drab and seem unfortunate.</i>	
Sadness	Contentment	Contentment	Excitement	Sadness	Sadness	
<i>It's beautiful. The colors are vibrant and inspiring.</i>	<i>Nice colors, feels life is starting.</i>	<i>The bright colors and the slant of the road bring about excitement in me, like a moving forward.</i>	<i>Industrial pollution provokes the anger.</i>	<i>The sad smog coming from the many chimneys of pollution is scary.</i>		
Awe	Excitement	Excitement	Anger	Sadness		
<i>Cute picture that makes me think of a warm summer day.</i>	<i>The light post is crooked making me think something hit it.</i>	<i>This place does not look so welcoming for some reason it looks spooky.</i>	<i>Destruction, abandoned, neglect, broken fences and lamps, disorganized, incomplete.</i>			
Contentment	Sadness	Fear	Disgust			

Fig. S8. More samples of the AIF dataset, where each anchor image has multiple corresponding text descriptions and each text description is categorized into the closest emotional category in Mikel's wheel [7].