

Image Composition Assessment with Saliency-augmented Multi-pattern Pooling

Bo Zhang
bo-zhang@sjtu.edu.cn

Li Niu*
ustcnewly@sjtu.edu.cn

Liqing Zhang
zhang-lq@cs.sjtu.edu.cn

MoE Key Lab of Artificial Intelligence
Shanghai Jiao Tong University
Shanghai, China

Abstract

Image composition assessment is crucial in aesthetic assessment, which aims to assess the overall composition quality of a given image. However, to the best of our knowledge, there is neither dataset nor method specifically proposed for this task. In this paper, we contribute the first composition assessment dataset CADB with composition scores for each image provided by multiple professional raters. Besides, we propose a composition assessment network SAMP-Net with a novel Saliency-Augmented Multi-pattern Pooling (SAMP) module, which analyses visual layout from the perspectives of multiple composition patterns. We also leverage composition-relevant attributes to further boost the performance, and extend Earth Mover’s Distance (EMD) loss to weighted EMD loss to eliminate the content bias. The experimental results show that our SAMP-Net can perform more favorably than previous aesthetic assessment approaches.

1 Introduction

Image aesthetic assessment aims to judge aesthetic quality automatically in a qualitative or quantitative way, which can be widely used in many down-stream applications such as assisted photo editing, intelligent photo album management, image cropping, and smartphone photography [8, 2, 11, 39, 40, 41, 43, 51]. Among the factors related to image aesthetics, image composition, which mainly concerns the arrangement of the visual elements inside the frame [53], is very critical in estimating image aesthetics [28, 36, 44], because composition directs the attention of viewer and has a significant impact on the aesthetic perception [12, 34, 38].

Despite the importance of image composition, there is no dataset readily available for image composition assessment. Some existing aesthetic datasets contain annotations related to image composition [8, 19, 22, 39]. However, they only have composition-relevant attributes without overall composition score except for PCCD dataset [9], but PCCD dataset only presents one reviewer’s composition rating for each image and this reviewer, an anonymous website visitor, may be unprofessional. So the ratings might be biased and inaccurate, which are far below the requirement for scientific evaluation. To this end, we contribute

*Corresponding author.

© 2021. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

a new image Composition Assessment DataBase (CADB) on the basis of Aesthetics and Attributes DataBase (AADB) dataset [22]. Our CADB dataset contains 9,497 images with each image rated by 5 individual raters who specialize in fine art for the overall composition quality. The details of our CADB dataset will be introduced in Section 3.

To the best of our knowledge, there is no method specifically designed for image composition assessment. However, some previous aesthetic assessment methods also take composition into consideration. We divide the existing composition-relevant approaches into two groups. 1) The composition-preserving methods [4, 32] can maintain image composition during both training and testing. However, these approaches fail to extract composition-relevant feature for composition assessment task. 2) The composition-aware approaches [28, 31, 52] extract composition-relevant feature by modeling the mutual dependencies between all pairs of objects or regions in the image. However, redundant and noisy information is likely to be introduced during this procedure, which may adversely affect the performance of composition assessment. Moreover, there are some previous methods [11, 10, 29, 49, 54, 55] designed to model the well-established photographic rules (*e.g.*, rule of thirds and golden ratio [20]), which humans use in evaluating image composition quality. However, these rule-based methods have two major limitations: 1) The hand-crafted feature extraction is tedious and laborious compared with deep learning features [27]. 2) Each rule is valid only for specific scenes and they did not consider which rules are applicable for a given scene [47].

Interestingly, composition pattern, as an important aspect of composition assessment, is not explicitly considered by the above methods. As shown in Figure 1, each composition pattern divides the holistic image into multiple non-overlapping partitions, which can model human perception of composition quality. In particular, by analyzing the visual layout (*e.g.*, positions and sizes of visual elements) according to composition pattern, *i.e.*, comparing the visual elements in various partitions, we can quantify the aesthetics of visual layout in terms of visual balance (*e.g.*, symmetrical balance and radial balance) [18, 23, 30], composition rules (*e.g.*, rule of thirds, diagonals and triangles) [24, 50], and so on. Different composition patterns offer different perspectives to evaluate composition quality. For example, the composition pattern in the top (*resp.*, bottom) row in Figure 1 can help judge the composition quality in terms of symmetrical (*resp.*, radial) balance.

To dissect visual layout based on different composition patterns, we propose a novel multi-pattern pooling module at the end of backbone to integrate the information extracted from multiple patterns, in which each pattern provides a perspective to evaluate the composition quality. Considering that the sizes and locations of salient objects are representative of visual layout and fundamental to image composition [30], we further integrate visual saliency [17] into our multi-pattern pooling module to encode the spatial and geometric information of salient objects, leading to our Saliency-Augmented Multi-pattern Pooling (SAMP) module. Additionally, since some composition patterns may play more important roles, we design weighted multi-pattern aggregation to fuse multi-pattern features, which can adaptively




Figure 1: Evaluating composition quality from the perspectives of different composition patterns. The first (*resp.*, second) row shows a good example and a bad example considering symmetrical (*resp.*, radial) balance.




Figure 2: The overall pipeline of our SAMP-Net for composition assessment. We use ResNet18 [14] as backbone. The detailed structure of our Saliency-Augmented Multi-pattern Pooling (SAMP) module and Attentional Attribute Feature Fusion (AAFF) module are illustrated in Figure 3 and Figure 4 respectively.

assign different weights to different patterns.

Moreover, because our dataset is built upon AADB dataset [22] with composition-relevant attributes, we further leverage composition-relevant attributes to boost the performance of composition assessment. Specifically, we propose an Attentional Attribute Feature Fusion (AAFF) module to fuse composition feature and attribute feature. Finally, after noticing the content bias existing in our dataset, that is, composition score distribution is severely influenced by object category, we extend Earth Mover’s Distance (EMD) loss in [15] to weighted EMD loss to eliminate the content bias.

The main contributions of this paper can be summarized as follows: 1) We contribute the first image composition assessment dataset CADB, in which each image has the composition scores annotated by five professional raters. 2) We propose a novel composition assessment method with Saliency-Augmented Multi-pattern Pooling (SAMP) module. 3) We investigate the effectiveness of auxiliary attributes and weighted EMD loss for composition assessment. 4) Our model outperforms previous aesthetic assessment methods on our dataset.

2 Related Work

2.1 Aesthetic Assessment Dataset

Many large-scale aesthetic assessment datasets have been collected in recent years, like Aesthetic Visual Analysis database (AVA) [55], AADB [22], Photo Critique Captioning Dataset (PCCD) [8], AVA-Comments [60], AVA-Reviews [53], FLICKER-AES [42], and DPC-Captions [49]. However, these datasets only have composition-relevant attributes without overall composition score, or only have one inaccurate composition score for each image, which are far below the requirement for composition assessment research. Unlike the existing aesthetic datasets, our CADB dataset contains composition ratings assigned to each image by multiple professional raters. Besides, we guarantee the reliability of our dataset based on sanity check and consistency analysis (see Section 3).

2.2 Composition-relevant Aesthetic Assessment

We can divide existing composition-relevant aesthetic assessment methods into traditional methods and deep learning methods. As surveyed in [2, 9], traditional methods [11, 25, 29, 33, 35, 36, 39, 42, 46, 49, 55, 58, 61] usually employed hand-crafted features or generic image features (*e.g.*, bag-of-visual-words [46] and Fisher vectors [67]) to learn image aesthetic evaluation, yet their generalization ability is limited by the complexity of image composition




Figure 3: Our designed eight composition patterns and Saliency-augmented Multi-pattern Pooling (SAMP) module.

assessment task. The deep learning based methods can be divided into two groups. The composition-preserving approaches [4, 5], without explicitly learning composition representations, produce inferior results on composition evaluation task. The composition-aware approaches [28, 31, 32] consider the relationship between all pairs of objects or regions in the image for modeling image composition, which is likely to introduce redundant and noisy information. Moreover, the above methods did not explicitly consider composition patterns. In contrast, we design a novel Saliency-Augmented Multi-pattern Pooling (SAMP) module, which provides an insightful and effective perspective for evaluating composition quality.

3 Composition Assessment DataBase (CADB)

To the best of our knowledge, there is no prior dataset specifically constructed for composition assessment. To support the research on this task, we build a dataset upon the existing AADB dataset [22], from which we collect a total of 9,958 real-world photos. We adopt a composition rating scale from 1 to 5, where a larger score indicates better composition. We make annotation guidelines for composition quality rating and train five individual raters who specialize in fine art. So for each image, we can obtain five composition scores ranging from 1 to 5. Given the subjective nature of human aesthetic activity [12, 33, 34], we perform sanity check and consistency analysis. Similar to [37], we use 240 additional ‘‘sanity check’’ images during annotating to roughly verify the validness of our annotations. We also examine the consistency of composition ratings provided by five individual raters (see Supplementary). Similar to [22, 35], we average the composition scores as the ground-truth composition mean score for each image, which is denoted as \bar{y} . More details about our CADB dataset will be elaborated in Supplementary.

Besides, we observe the content bias in our CADB dataset, that is, there are some biased categories whose score distributions are concentrated in a very narrow interval. After removing 461 biased images, we split the remaining images into 8,547 training images and 950 test

images, in which the test set is made less biased for better evaluation (see Supplementary).

4 Methodology

To accomplish the composition assessment task, we propose a novel network SAMP-Net, which is named after Saliency-Augmented Multi-pattern Pooling (SAMP) module. The overall pipeline of our method is illustrated in Figure 2, where we first extract the global feature map from input image by backbone (*e.g.*, ResNet18 [14]) and then yield aggregated pattern feature through our SAMP module, which is followed by Attentional Attribute Feature Fusion (AAFF) module to fuse the composition feature and attribute feature. After that, we predict composition score distribution based on the fused feature and predict the attribute score based on the attribute feature, which are supervised by weighted EMD loss and Mean Squared Error (MSE) loss respectively.

4.1 Saliency-augmented Multi-pattern Pooling

Multi-pattern Pooling: As demonstrated in Figure 3(a), we empirically design eight basic composition patterns inspired by classic composition guidelines. For instance, pattern 1,2,6,7 are inspired by symmetrical composition. Pattern 3,4 are inspired by diagonal composition. Pattern 5 is inspired by centre composition. Pattern 8 is inspired by rule of thirds [24, 50]. Although our pattern design is inspired by composition rules, there is no strict one-to-one correspondence between composition rules and patterns. Each pattern provides a perspective for evaluating composition quality, which may be beyond the scope of a single rule. For example, Pattern 8 is related to rule of thirds, but not limited to rule of thirds. Based on Pattern 8, more useful information can be excavated by comparing the visual elements in nine partitions.

Since humans typically employ multiple perspectives when analysing image composition, composition assessment should be accomplished based on all composition patterns in a comprehensive way. Therefore, we propose multi-pattern pooling to achieve this goal, which is illustrated in Figure 3(b). Given an $H \times W$ global feature map \mathbf{F} with C channels, which is extracted from input image by backbone, we represent the pixel-wise feature at each location as $\mathbf{x}_{i,j}$, where $0 < i \leq H$, $0 < j \leq W$. For the p -th pattern, we divide \mathbf{F} into K_p non-overlapping partitions $\{\mathcal{X}_1^p, \mathcal{X}_2^p, \dots, \mathcal{X}_{K_p}^p\}$ and K_p is the total number of partitions in this pattern. Then, the feature of the k -th partition can be obtained via average pooling: $\theta(\mathcal{X}_k^p) = \frac{1}{|\mathcal{X}_k^p|} \sum_{(i,j) \in \mathcal{X}_k^p} \mathbf{x}_{i,j} \in \mathbb{R}^C$.

Saliency-augmented Multi-pattern Pooling: Considering the significance of salient objects for composition assessment, we further incorporate the saliency information (*i.e.*, locations and scales of salient objects) into multi-pattern pooling. To achieve this goal, we utilize an unsupervised saliency detection method [17] to produce saliency maps for input images. We have also tried several supervised methods [6, 16, 59], which prove to be less effective. After obtaining the saliency map, we downsample it to $H_{sal} \times W_{sal}$ through max pooling. Recall that the size of global feature map is $H \times W$, we set $H_{sal} = 8H$ and $W_{sal} = 8W$ for retaining more details of salient objects.

Different from $\theta(\mathcal{X}_k^p)$ using average pooling, we directly reshape each partition of saliency map into a vector, because the pooling operation will result in significant information loss. Specifically, for the k -th partition in the p -th pattern, we reshape the saliency map in this

partition into a saliency vector $\psi(\mathcal{X}_k^p) \in \mathbb{R}^{D_k^p}$, in which D_k^p varies with partition and pattern. Then, we concatenate $\psi(\mathcal{X}_k^p)$ and $\theta(\mathcal{X}_k^p)$ to generate the partition feature $[\psi(\mathcal{X}_k^p), \theta(\mathcal{X}_k^p)]$.

For the p -th pattern, we concatenate the partition features of K_p partitions into a long vector $\tilde{\mathbf{f}}_{samp}^p$, which is followed by a fc layer and ReLU activation function to produce the pattern vector $\mathbf{f}_{samp}^p \in \mathbb{R}^{C'}$. Intuitively, $[\psi(\mathcal{X}_k^p), \theta(\mathcal{X}_k^p)]$ extracts the visual information in each partition and \mathbf{f}_{samp}^p encodes the relationship among visual elements in different partitions.

Weighted Multi-pattern Aggregation: Since some composition patterns may play more important roles when evaluating image composition, our model is trained to assign different weights for different patterns. Precisely, we apply global average pooling, a fc layer, and softmax normalization to the global feature map \mathbf{F} , producing the multi-pattern weight w_p for the p -th pattern. Then, we have the aggregated pattern feature via weighted summation $\mathbf{f}_{samp} = \sum_{p=1}^P w_p \mathbf{f}_{samp}^p$, in which P is the number of composition patterns ($P = 8$). Based on the learnt weights, we can know the dominant patterns in determining the overall composition quality and provide interpretable guidance for users (see Section 5.4).

Comparison with Spatial Pyramid Pooling: Although the proposed SAMP and Spatial Pyramid Pooling (SPP) [13] are similar in architecture, both of which pool features from multiple sets of partitions, SAMP is significantly different from SPP in three main aspects: 1) our pooling patterns are specifically designed and well-tailored for image composition evaluation, which can analyse the composition quality from the viewpoint of composition patterns; 2) we introduce visual saliency into multi-pattern pooling; 3) we learn pattern weights which provide interpretable guidance for improving composition quality.

4.2 Attentional Attribute Feature Fusion

Since our dataset is built upon AADB [22], which is associated with composition-relevant attributes, it is natural to consider using them to help composition assessment. We use five composition-relevant attribute annotations: *rule of thirds*, *balancing elements*, *object emphasis*, *symmetry*, and *repetition*.

Specifically, as illustrated in Figure 2, we decompose the aggregated pattern feature $\mathbf{f}_{samp} \in \mathbb{R}^{C'}$ into composition feature \mathbf{f}_{comp} and attribute feature \mathbf{f}_{atts} by using two separate fc layers, the dimensions of which are both set to $\frac{C'}{2}$. We dynamically weigh the contributions of \mathbf{f}_{comp} and \mathbf{f}_{atts} for the composition assessment task, as illustrated in Figure 4. First, we apply a fc layer and sigmoid activation to the concatenation of \mathbf{f}_{comp} and \mathbf{f}_{atts} , to learn the attention coefficients $[e_1, e_2]$ for two types of features. Then, we concatenate the weighted composition feature and attribute feature, yielding the fused feature $\mathbf{f}_{fused} = [e_1 \mathbf{f}_{comp}, e_2 \mathbf{f}_{atts}] \in \mathbb{R}^{C'}$. During training, an additional layer is added to perform attribute prediction based on the attribute feature \mathbf{f}_{atts} . We employ MSE loss denoted as \mathcal{L}_{atts} for attribute prediction.

As mentioned in Section 3, we observe the content bias in our dataset, in which case the network may find a shortcut to simply rate images based on their contents. To mitigate the content bias in training set, we extend EMD loss to weighted EMD loss denoted as \mathcal{L}_{wEMD}




Figure 4: Attentional Attribute Feature Fusion (AAFF) module. fc means a fully-connected layer with sigmoid activation and e_1, e_2 are attention coefficients.

	<i>WE</i>	<i>MP</i>	<i>PW</i>	<i>SA</i>	<i>AF</i>	<i>AA</i>	MSE↓	EMD↓	SRCC↑	LCC↑
1							0.4534	0.1943	0.6025	0.6148
2	✓						0.4373	0.1859	0.6105	0.6258
3	✓	✓					0.4170	0.1847	0.6292	0.6435
4	✓	✓	✓				0.4134	0.1829	0.6323	0.6483
5	✓	✓	✓	✓			0.4088	0.1820	0.6421	0.6544
6	✓	†	✓	✓			0.4274	0.1854	0.6226	0.6293
7	✓	‡	✓	✓			0.4205	0.1845	0.6319	0.6363
8	✓				✓		0.4320	0.1850	0.6200	0.6303
9	✓	✓	✓	✓	✓		0.3979	0.1817	0.6439	0.6610
10	✓	✓	✓	✓	✓	✓	0.3867	0.1798	0.6564	0.6709

Table 1: Ablation studies of different components in our model. † means Spatial Pyramid Pooling (SPP) [14]. ‡ means Multi-scale Pyramid Pooling (MPP) [66]. *WE* means weighted EMD loss. *MP* means multi-pattern pooling. *PW* means pattern weights. *SA* means saliency-augmented. *AF* indicates attribute feature and *AA* indicates attentional attribute feature fusion.

(see Supplementary), which assigns smaller weights to biased samples when calculating EMD Loss. Finally, our SAMP-Net can be trained in an end-to-end manner with attribute prediction loss \mathcal{L}_{atts} and weighted EMD loss \mathcal{L}_{wEMD} :

$$\mathcal{L} = \mathcal{L}_{wEMD} + \lambda \mathcal{L}_{atts}, \quad (1)$$

where λ is a trade-off parameter set as 0.1 via cross validation.

5 Experiments

5.1 Implementation Details and Evaluation Metric

We use ResNet18 [14] pretrained on ImageNet [8] as the backbone of our SAMP-Net. Unless otherwise specified, all input images are resized to 224×224 for both training and testing following [21, 26, 45], leading to a global feature map of $H \times W = 7 \times 7$, and the saliency map is downsampled to $H_{sal} \times W_{sal} = 56 \times 56$ before passing to the SAMP. More details can be found in Supplementary. All experiments are conducted on our CADB dataset.

To evaluate the composition score distribution and composition mean score predicted by different models, it is natural to adopt EMD and MSE as the evaluation metrics. EMD measures the closeness between the predicted and ground-truth composition score distributions as in [15]. MSE is computed between the predicted and ground-truth composition mean scores. Moreover, following existing aesthetic assessment approaches [9, 22, 48], we also report the ranking correlation measured by Spearman’s Rank Correlation Coefficient (SRCC) and the linear association measured by Linear Correlation Coefficient (LCC) between the predicted and ground-truth composition mean scores.

5.2 Ablation Study

To evaluate the effectiveness of each individual component in our SAMP-Net, we conduct a series of experiments and report all the evaluation metrics described in Section 5.1. In this section, we start from ResNet18 backbone and build up our holistic model step by step.

Method	MSE↓	EMD↓	SRCC↑	LCC↑
ResNet18	0.4534	0.1943	0.6025	0.6148
AADB [22]	0.4234	0.1923	0.6236	0.6415
MNA-CNN [32]	0.4260	0.1944	0.6108	0.6375
A-Lamp [31]	0.4230	0.1898	0.6270	0.6456
VP-Net [24]	0.4304	0.1948	0.6169	0.6285
RG-Net [28]	0.4398	0.1915	0.6026	0.6218
AFDC-Net [9]	0.4245	0.1910	0.6154	0.6388
SAMP-Net (Ours)	0.3867	0.1798	0.6564	0.6709

Table 2: Comparison of different methods on the composition assessment task. All models are trained and evaluated on the proposed CADB dataset.

Weighted EMD Loss: We start from basic ResNet18 [14], and report the results using EMD loss and weighted EMD loss in Table 1. The experimental results show that training with weighted EMD loss (row 2) performs better than standard EMD loss (row 1) with a clear gap of test EMD between these two models, which is attributed to the advantage of weighted EMD loss in eliminating content bias.

Saliency-Augmented Multi-pattern Pooling (SAMP): Based on ResNet18 with weighted EMD loss (row 2), we add our SAMP module and also explore its ablated versions. We first investigate vanilla multi-pattern pooling without saliency or pattern weights (row 3), in which saliency vector is excluded from partition feature and the pattern features of multiple patterns are simply averaged. Then, we learn pattern weights to aggregate multiple pattern features (row 4). By comparing row 3 and row 4, it is beneficial to adaptively assign different weights to different pattern features. We further incorporate saliency map into SAMP module (row 5). The comparison between row 4 and row 5 proves that is useful to emphasize the layout information of salient objects. Considering the architecture similarity between Spatial Pyramid Pooling (SPP) [13] and our multi-pattern pooling, we replace our multi-pattern pooling with SPP using scales $\{1 \times 1, 2 \times 2, 3 \times 3\}$ following [9] (row 6). In addition, we also show the results of using Multi-scale Pyramid Pooling (MPP) [56] in row 7, in which we make an image pyramid containing three scaled images. The comparisons (row 5 v.s. row 6, row 5 v.s. row 7) show that the model using multi-pattern pooling outperforms both SPP and MPP. The reason is that our proposed multi-pattern pooling is specifically designed and well-tailored for composition assessment task.

Attentional Attribute Feature Fusion (AAFF): Built on row 2 (*resp.*, row 5) in Table 1, we additionally learn attribute feature and directly concatenate it with composition feature, leading to row 8 (*resp.*, row 9). The experimental results demonstrate that composition-relevant attributes can help boost the performance of composition evaluation. This sheds light on that composition-relevant attribute prediction and composition evaluation are two related and reciprocal tasks. Finally, we complete our attentional attribute feature fusion module by learning weights for weighted concatenation (row 10). From row 9 and row 10, we can observe that the model using weighted concatenation is better than that using plain concatenation, which validates the superiority of attentional fusion mechanism.

5.3 Comparison with Existing Methods

To the best of our knowledge, there is no method specifically designed for image composition assessment. Nevertheless, some previous aesthetic assessment methods [9, 22, 28, 31, 32,





		<p>Ground-truth Score: 1.80</p> <p>Predicted Score: 2.48</p> <p>Pattern Weights: [0.25, 0.13, 0.08, 0.08, 0.07, 0.05, 0.13, 0.21]</p>			<p>Ground-truth Score: 2.40</p> <p>Predicted Score: 2.74</p> <p>Pattern Weights: [0.10, 0.27, 0.01, 0.05, 0.14, 0.15, 0.05, 0.22]</p>
		<p>Ground-truth Score: 4.00</p> <p>Predicted Score: 3.96</p> <p>Pattern Weights: [0.25, 0.06, 0.35, 0.02, 0.19, 0.06, 0.02, 0.04]</p>			<p>Ground-truth Score: 3.00</p> <p>Predicted Score: 2.66</p> <p>Pattern Weights: [0.17, 0.20, 0.01, 0.24, 0.04, 0.08, 0.07, 0.19]</p>
		<p>Ground-truth Score: 1.80</p> <p>Predicted Score: 1.83</p> <p>Pattern Weights: [0.07, 0.09, 0.03, 0.06, 0.65, 0.06, 0.00, 0.04]</p>			<p>Ground-truth Score: 3.00</p> <p>Predicted Score: 3.06</p> <p>Pattern Weights: [0.22, 0.09, 0.02, 0.03, 0.11, 0.25, 0.06, 0.21]</p>
		<p>Ground-truth Score: 1.60</p> <p>Predicted Score: 1.93</p> <p>Pattern Weights: [0.15, 0.12, 0.03, 0.15, 0.05, 0.06, 0.24, 0.19]</p>			<p>Ground-truth Score: 4.00</p> <p>Predicted Score: 4.02</p> <p>Pattern Weights: [0.12, 0.37, 0.00, 0.00, 0.04, 0.01, 0.02, 0.44]</p>
Image & Dominant Pattern	Saliency Map	Ground-truth and Predicted Composition Mean Score & Predicted Pattern Weights	Image & Dominant Pattern	Saliency Map	Ground-truth and Predicted Composition Mean Score & Predicted Pattern Weights

Figure 5: Analysis of the correlation between an image and its dominant pattern with the largest weight. We show the estimated pattern weights and the largest weight is colored green. We also show the ground-truth/predicted composition mean score in blue/red.

explicitly take composition into consideration. Since most of these methods do not yield score distribution, we make a slight modification on the prediction layer of these methods to be compatible with EMD loss [15]. For fair comparison, all methods are trained and tested on our CADB dataset with ResNet18 pretrained on ImageNet [8] as backbone.

In Table 2, we compare our method with different composition-relevant aesthetic assessment methods. The baseline model (ResNet18) only consists of the pretrained ResNet18 and a prediction head, which is the same as row 1 in Table 1. Among these baselines, A-Lamp is the most competitive one, probably because A-Lamp introduces additional saliency information to learn the pairwise spatial relationship between objects. Our SAMP-Net clearly outperforms all the composition-relevant baselines, which demonstrates that our method is more adept at image composition assessment.

5.4 Analysis of Composition Pattern

To take a close look at the learnt pattern weights which indicate the importance of different patterns on the overall composition quality, we show the input image, its saliency map, its ground-truth/predicted composition mean score, and its pattern weights in Figure 5.

For each image, the composition pattern with the largest weight is referred to as its dominant pattern. For each pattern, we show one example image with this pattern as dominant pattern and overlay this pattern on the image in Figure 5, which reveals from which perspective the input image is given a high or low score. For example, in the right figure of the last row, the surfer is placed at the intersection point between the gridlines of pattern 8, which implicates that the image conforms to the rule of thirds properly, yielding a relatively high score. On the contrary, in the right figure of the first row, the arch slightly deviates from its



Figure 6: We show some failure cases in the test set, which have the highest absolute errors between the predicted composition mean scores (out of bracket) and the ground-truth composition mean scores (in bracket).

symmetrical axis under pattern 2. So the low score implies that maintaining horizontal symmetry may enhance the composition quality. In the left figure of the third row, per the low score under pattern 5, the dog is suggested to be moved to the center. In summary, our SAMP module can facilitate composition assessment by integrating the information from multiple patterns and provide constructive suggestions for improving the composition quality.

5.5 Additional Experiments in Supplementary

Due to the space limitation, we present some experiments in Supplementary, including the results of using different training set sizes, backbones, and hyper-parameters λ in (1), weighted EMD loss analysis, the effectiveness of each pattern, the impact of using more composition patterns, comparison with the performance of human raters, more results on the CADB and PCCD [9] datasets.

5.6 Limitations

While our method can generally achieve accurate and reliable composition assessment, it still has some failure cases. We show several failure cases in Figure 6, which have the highest absolute errors between the predicted and ground-truth composition mean scores. We can observe that our model tends to predict relatively low scores for these images with high composition mean scores, which is probably due to the distracting backgrounds and complicated composition patterns. In addition, there is a clear gap between our method and human raters on ranking the composition quality of different images (see Supplementary), which needs to be enhanced in the future work.

6 Conclusion

In this paper, we have contributed the first composition assessment dataset CADB with five composition scores for each image. We have also proposed a novel method SAMP-Net with saliency-augmented multi-pattern pooling. Equipped with SAMP module, AAFF module, and weighted EMD loss, our method is capable of achieving the best performance for composition assessment.

Acknowledgement

This work is sponsored by National Natural Science Foundation of China (Grant No. 61902247) and Shanghai Sailing Program (19YF1424400).

References

- [1] S. Bhattacharya, R. Sukthankar, and M. Shah. A framework for photo-quality assessment and enhancement based on visual aesthetics. In *ACM-Multimedia*, 2010.
- [2] A. Brachmann and C. Redies. Computational and experimental approaches to visual aesthetics. *Frontiers in computational neuroscience*, 11(1):102–119, 2017.
- [3] K. Chang, KH. Lu, and CS. Chen. Aesthetic critiques generation for photos. In *ICCV*, 2017.
- [4] Q. Chen, W. Zhang, N. Zhou, P. Lei, Y. Xu, Y. Zheng, and J. Fan. Adaptive fractional dilated convolution network for image aesthetics assessment. In *CVPR*, 2020.
- [5] Y. Chen, J. Klopp, M. Sun, S. Chien, and K. Ma. Learning to compose with professional photographs on the web. In *ACM-Multimedia*, 2017.
- [6] M. Cornia, L. Baraldi, G. Serra, and R. Cucchiara. Predicting human eye fixations via an LSTM-based saliency attentive model. *IEEE Transactions on Image Processing*, 27(10):5142–5154, 2018.
- [7] R. Datta, D. Joshi, J. Li, and J. Wang. Studying aesthetics in photographic images using a computational approach. In *ECCV*, 2006.
- [8] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [9] Y. Deng, C.C. Loy, and X. Tang. Image aesthetic assessment: An experimental survey. *IEEE Signal Processing Magazine*, 34(4):80–106, 2017.
- [10] S. Dhar, V. Ordonez, and T. Berg. High level describable attributes for predicting aesthetics and interestingness. In *CVPR*, 2011.
- [11] Yuming Fang, Hanwei Zhu, Yan Zeng, Kede Ma, and Zhou Wang. Perceptual quality assessment of smartphone photography. In *CVPR*, 2020.
- [12] M. Freeman. *The photographer’s eye: Composition and design for better digital photos*. CRC Press, 2007.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1904–1916, 2015.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [15] L. Hou, C.P. Yu, and D. Samaras. Squared earth mover’s distance-based loss for training deep neural networks. *ArXiv*, abs/1611.05916, 2016.
- [16] Q. Hou, MM. Cheng, X. Hu, A. Borji, Z. Tu, and P. Torr. Deeply supervised salient object detection with short connections. In *CVPR*, 2017.
- [17] X. Hou and L. Zhang. Saliency detection: A spectral residual approach. In *CVPR*, 2007.

- [18] A. Jahanian, S. Vishwanathan, and J. Allebach. Learning visual balance from large-scale datasets of aesthetically highly rated images. In *Human Vision and Electronic Imaging XX*, 2015.
- [19] X. Jin, L. Wu, G. Zhao, X. Li, X. Zhang, S. Ge, D. Zou, B. Zhou, and X. Zhou. Aesthetic attributes assessment of images. In *ACM-Multimedia*, 2019.
- [20] 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.
- [21] Keunsoo Ko, Jun-Tae Lee, and Chang-Su Kim. PAC-Net: Pairwise aesthetic comparison network for image aesthetic assessment. In *ICIP*, 2018.
- [22] S. Kong, X. Shen, Z. Lin, R. Mech, and C. Fowlkes. Photo aesthetics ranking network with attributes and content adaptation. In *ECCV*, 2016.
- [23] J.T. Lee, H. Kim, C. Lee, and C. Kim. Semantic line detection and its applications. In *ICCV*, 2017.
- [24] J.T. Lee, H. Kim, C. Lee, and C. Kim. Photographic composition classification and dominant geometric element detection for outdoor scenes. *Journal of Visual Communication and Image Representation*, 55(1):91–105, 2018.
- [25] C. Li, A. Gallagher, A. Loui, and T. Chen. Aesthetic quality assessment of consumer photos with faces. In *ICIP*, 2010.
- [26] Leida Li, Hancheng Zhu, Sicheng Zhao, Guiguang Ding, and Weisi Lin. Personality-assisted multi-task learning for generic and personalized image aesthetics assessment. *IEEE Transactions on Image Processing*, 29(1):3898–3910, 2020.
- [27] Xuewei Li, Xueming Li, Gang Zhang, and Xianlin Zhang. A novel feature fusion method for computing image aesthetic quality. *IEEE access*, 8:63043–63054, 2020.
- [28] D. Liu, R. Puri, N. Kamath, and S. Bhattacharya. Composition-aware image aesthetics assessment. In *WACV*, 2020.
- [29] Ligang Liu, Renjie Chen, Lior Wolf, and Daniel Cohen-Or. Optimizing photo composition. In *Computer Graphics Forum*, 2010.
- [30] S. Lok, S. Feiner, and G. Ngai. Evaluation of visual balance for automated layout. In *Proceedings of the 9th international conference on Intelligent user interfaces*, 2004.
- [31] S. Ma, J. Liu, and C. Chen. A-Lamp: Adaptive layout-aware multi-patch deep convolutional neural network for photo aesthetic assessment. In *CVPR*, 2017.
- [32] L. Mai, H. Jin, and F. Liu. Composition-preserving deep photo aesthetics assessment. In *CVPR*, 2016.
- [33] L. Marchesotti, F. Perronnin, D. Larlus, and G. Csurka. Assessing the aesthetic quality of photographs using generic image descriptors. In *ICCV*, 2011.
- [34] B. Martinez and J. Block. *Visual forces: an introduction to design*. Pearson College Division, 1995.

- [35] N. Murray, L. Marchesotti, and F. Perronnin. AVA: A large-scale database for aesthetic visual analysis. In *CVPR*, 2012.
- [36] P. Obrador, L. Schmidt-Hackenberg, and N. Oliver. The role of image composition in image aesthetics. In *ICIP*, 2010.
- [37] F. Perronnin and C. Dance. Fisher kernels on visual vocabularies for image categorization. In *CVPR*, 2007.
- [38] D. Präkel. *The fundamentals of creative photography*. Bloomsbury Publishing, 2010.
- [39] Yogesh Singh Rawat and Mohan S Kankanhalli. Context-aware photography learning for smart mobile devices. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 12(1):1–24, 2015.
- [40] Yogesh Singh Rawat and Mohan S Kankanhalli. Clicksmart: A context-aware viewpoint recommendation system for mobile photography. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(1):149–158, 2016.
- [41] Yogesh Singh Rawat, Mingli Song, and Mohan S Kankanhalli. A spring-electric graph model for socialized group photography. *IEEE Transactions on Multimedia*, 20(3):754–766, 2017.
- [42] J. Ren, X. Shen, Z. Lin, R. Mech, and D. Foran. Personalized image aesthetics. In *ICCV*, 2017.
- [43] R. Sukthankar S. Bhattacharya and M. Shah. A holistic approach to aesthetic enhancement of photographs. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 7(1):1–21, 2011.
- [44] A. Savakis, S. Etz, and A. Loui. Evaluation of image appeal in consumer photography. In *Human vision and electronic imaging V*, 2000.
- [45] Katharina Schwarz, Patrick Wieschollek, and Hendrik PA Lensch. Will people like your image? learning the aesthetic space. In *WACV*, 2018.
- [46] H. Su, T. Chen, C. Kao, W. Hsu, and S. Chien. Scenic photo quality assessment with bag of aesthetics-preserving features. In *ACM-Multimedia*, 2011.
- [47] Yu-Chuan Su, Raviteja Vemulapalli, Ben Weiss, Chun-Te Chu, Philip Andrew Mansfield, Lior Shapira, and Colvin Pitts. Camera view adjustment prediction for improving image composition. *arXiv preprint arXiv:2104.07608*, 2021.
- [48] H. Talebi and P. Milanfar. NIMA: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):3998–4011, 2018.
- [49] X. Tang, W. Luo, and X. Wang. Content-based photo quality assessment. *IEEE Transactions on Image Processing*, 15(8):1930–1943, 2013.
- [50] K. Thömmes and R. Hübner. Instagram likes for architectural photos can be predicted by quantitative balance measures and curvature. *Frontiers in psychology*, 9(1):1050–1067, 2018.

- [51] Yi Tu, Li Niu, Weijie Zhao, Dawei Cheng, and Liqing Zhang. Image cropping with composition and saliency aware aesthetic score map. In *AAAI*, 2020.
- [52] W. Wang and R. Deng. Modeling human perception for image aesthetic assessment. In *ICIP*, 2019.
- [53] W. Wang, S. Yang, W. Zhang, and J. Zhang. Neural aesthetic image reviewer. *IET Computer Vision*, 13(8):749–758, 2019.
- [54] Min-Tzu Wu, Tse-Yu Pan, Wan-Lun Tsai, Hsu-Chan Kuo, and Min-Chun Hu. High-level semantic photographic composition analysis and understanding with deep neural networks. In *ICMEW*, 2017.
- [55] Yaowen Wu, Christian Bauckhage, and Christian Thurau. The good, the bad, and the ugly: Predicting aesthetic image labels. In *ICPR*, 2010.
- [56] Donggeun Yoo, Sunggyun Park, Joon-Young Lee, and In So Kweon. Multi-scale pyramid pooling for deep convolutional representation. In *CVPRW*, 2015.
- [57] N. Yu, X. Shen, L. Lin, R. Mech, and C. Barnes. Learning to detect multiple photographic defects. In *WACV*, 2018.
- [58] L. Zhang, Y. Gao, R. Zimmermann, Q. Tian, and X. Li. Fusion of multichannel local and global structural cues for photo aesthetics evaluation. *IEEE Transactions on Image Processing*, 23(3):1419–1429, 2014.
- [59] T. Zhao and X. Wu. Pyramid feature attention network for saliency detection. In *CVPR*, 2019.
- [60] Y. Zhou, X. Lu, J. Zhang, and J.Z. Wang. Joint image and text representation for aesthetics analysis. In *ACM-Multimedia*, 2016.
- [61] Z. Zhou, S. He, J. Li, and J.Z. Wang. Modeling perspective effects in photographic composition. In *ACM-Multimedia*, 2015.