

Image Cropping with Composition and Saliency Aware Aesthetic Score Map

Yi Tu,¹ Li Niu,^{1*} Weijie Zhao,² Dawei Cheng,¹ Liqing Zhang^{1*}

¹MoE Key Lab of Artificial Intelligence, Department of Computer Science and Engineering
Shanghai Jiao Tong University, Shanghai, China
{tuyi1991, ustcnewly, dawei.cheng}@sjtu.edu.cn, zhang-lq@cs.sjtu.edu.cn

²Versa Inc, Shanghai, China
weijie.zhao@versa-ai.com

Abstract

Aesthetic image cropping is a practical but challenging task which aims at finding the best crops with the highest aesthetic quality in an image. Recently, many deep learning methods have been proposed to address this problem, but they did not reveal the intrinsic mechanism of aesthetic evaluation. In this paper, we propose an interpretable image cropping model to unveil the mystery. For each image, we use a fully convolutional network to produce an aesthetic score map, which is shared among all candidate crops during crop-level aesthetic evaluation. Then, we require the aesthetic score map to be both composition-aware and saliency-aware. In particular, the same region is assigned with different aesthetic scores based on its relative positions in different crops. Moreover, a visually salient region is supposed to have more sensitive aesthetic scores so that our network can learn to place salient objects at more proper positions. Such an aesthetic score map can be used to localize aesthetically important regions in an image, which sheds light on the composition rules learned by our model. We show the competitive performance of our model in the image cropping task on several benchmark datasets, and also demonstrate its generality in real-world applications.

1 Introduction

Given an image, the image cropping task aims at finding the crops with the best aesthetic quality. It is an important task that can be widely used in a lot of down-stream applications, *e.g.*, photo post-processing (Chen et al. 2017b), view recommendation (Li et al. 2018; Wei et al. 2018), and image thumbnailing (Esmaeili, Singh, and Davis 2017). In order to find the best crop, an image cropping model will first generate a large number of candidate crops and then determine the best crop based on crop-level aesthetic evaluation. So an image cropping model is usually composed of two stages, candidate generation and aesthetic evaluation. A good image crop is achieved by selecting important contents and placing them with a good composition. The required knowledge for such a task can be categorized into two parts, *i.e.*, content preference and composition preference. Therefore, a good




Figure 1: Images crop with composition rules. The orange box in each image denotes a good crop found based on human-defined composition rules. The white dotted lines denote the auxiliary lines used in these composition rules.

image cropping model should be able to learn and leverage such preferences when searching for the best crop.

Early methods achieve this goal by explicitly utilizing some photography knowledge like human-defined composition rules, *e.g.*, Rule of Thirds and Rule of Central (See Figure 1). With the development of deep learning, recent researchers learn image cropping in a data-driven manner and many aesthetic datasets are constructed to encode the aesthetic preference of humans. Recent methods (Wang and Shen 2017; Chen et al. 2017b; Wei et al. 2018; Lu et al. 2019c) treat it as an object detection task. They used aesthetic datasets to train an aesthetic evaluation model and applied it to compare candidate crops. Due to the power of deep learning, these methods have brought progresses in this field, but the intrinsic mechanism remains unrevealed.

In this paper, we propose an interpretable image cropping model to produce both composition-aware and saliency-aware Aesthetic Score Maps, called ASM-Net. Our approach was first inspired by the Class-Activation-Map (CAM) method (Zhou et al. 2016), which uses a class activation map to localize the most discriminative image regions in image classification task. Similarly, we expect to use an aesthetic score map to localize aesthetically important image regions. The aesthetic score of a region can be obtained via average pooling and the regions with larger aesthetic scores are of higher aesthetic quality. However, direct application of CAM has been proven ineffective because the aesthetic evaluation task is more complicated than classification and one region cannot be simply represented by a single score.

*Corresponding author.

For example, in Rule of Central (left panel in Figure 1), we always place the key object at the center to achieve a symmetric and balanced photo crop. So the aesthetic score of the key object located at the center should be larger than off-center. So we realize that the aesthetic score of a region should vary with its relative position in different crops.

To achieve this goal, we first design some composition patterns to accommodate composition rules so that our model is capable of capturing different composition preferences in the training data. Each composition pattern divides a crop into multiple non-overlapping composition partitions (see Figure 3). Given a composition pattern, the aesthetic score of a region depends on the composition partition it belongs to in different crops. We name this property as “composition-aware”. In this case, one region is not represented by a single score, but a set of composition-aware scores. Based on such a composition-aware aesthetic score map, we design composition-aware pooling to calculate crop-level aesthetic scores efficiently.

Furthermore, we believe that our aesthetic score map should focus more on visually salient objects, so we introduce visual saliency into our model as extra supervision. Previous saliency-based image cropping methods (Wang and Shen 2017; Lu et al. 2019b) generally assume that the most salient object is the most important content and must be included in the best crop. However, we find that this assumption does not hold in many real-world images. For example, landscape photos could have no salient objects, while a party photo might have multiple salient figures that are equally important. These methods could fail in such images.

Unlike previous saliency-based methods, our assumption is more flexible and practical: a salient region should be more sensitive to the composition partition it belongs to than other regions. We design a novel saliency-aware loss to realize this assumption. By enforcing salient regions to be more composition-sensitive, our model could learn to place salient objects in the proper positions. With such a composition-aware and saliency-aware aesthetic score map, we could unveil the intrinsic mechanism of image cropping by locating the aesthetically important regions. Moreover, our model can be easily applied to a wide range of real-world applications. The main contributions of our work are as follows:

- We propose a novel image cropping model ASM-Net with composition-aware and saliency-aware aesthetic score map, which can encode content preference and composition preference.
- Our model is able to unveil the intrinsic mechanism of image cropping, by localizing aesthetically important regions based on our aesthetic score map.
- Our model outperforms the state-of-the-art methods on three benchmark datasets and has good generality to real-world applications.

2 Related Work

Aesthetic Evaluation: Recently, some large scale aesthetic datasets, like AVA (Murray, Marchesotti, and Perronnin 2012) and AADB (Kong et al. 2016), have been constructed, which enables learning aesthetic evaluation models in a

data-driven manner. By using these datasets, many deep learning methods (Lu et al. 2014; 2015; Hosu, Goldlucke, and Saupé 2019) have been proposed for image-level aesthetic evaluation. However, some researchers realized that we need more fine-grained aesthetic annotations for image cropping so that many similar crops could be compared. Thus, Chen et al. (2017a) first proposed a dataset to compare randomly sampled crops from one image. Following this idea, Wei et al. (2018) increased the number of crops in each image and proposed the first densely annotated image cropping dataset, in which each image has 24 annotated crops with aesthetic scores. Zeng et al. (2019) took a further step and presented another densely annotated dataset with about 85 crops for each image. In our method, we use such datasets to train our image cropping model.

Aesthetic-based Image Cropping: Aesthetic-based image cropping improves the cropping results by increasing their aesthetic quality. Early methods (Zhang et al. 2013; Yan et al. 2013; Fang et al. 2014) achieve this goal by using hand-crafted features and composition rules. By introducing deep learning into this area, some methods proposed to solve image cropping in a data-driven manner. Wei et al. (2018) used an object detection framework (Liu et al. 2016) with two networks, one for generating candidate crops and one for aesthetic evaluation. More recently, (Zeng et al. 2019) proposed new metrics to evaluate the performance of image cropping models. Compared with these deep learning-based approaches, our method is more interpretable.

Saliency-based Image Cropping: The saliency-based methods (Sun and Ling 2013; Chen et al. 2016; Wang and Shen 2017; Lu et al. 2019b; 2019a; Li et al. 2019) focused on preserving the most important content in the best crop. Wang and Shen (2017) used a saliency dataset (Jiang et al. 2015) to train a network to generate candidate crops that cover the most salient region. Similarly, Lu et al. (2019b) proposed to generate an initial visual saliency rectangle to include the most important objects. Unlike these methods, our method has a different assumption on the relationship between visual saliency and aesthetic evaluation.

3 Methodology

The flowchart of our method is illustrated in Figure 2.

3.1 Composition-Aware Aesthetic Score Map

Aesthetic Score Map: Similar to Class-Activation-Map (CAM) (Zhou et al. 2016), which can localize discriminative regions based on class activation map, we attempt to obtain an aesthetic score map for ease of localizing aesthetically important regions, which can be achieved by Fully Convolutional Network (Chen et al. 2015; Long, Shelhamer, and Darrell 2015). Considering that a high-resolution score map could be beneficial for region localization and imaging cropping, we upsample the feature maps from different layers to the input image size and concatenate them as our feature map. Then, we apply a 1×1 convolutional layer to transform this feature map to an aesthetic score map, with a larger score indicating higher aesthetic quality.

We use $\mathcal{I} = \{p_{i,j} | 0 < i \leq H, 0 < j \leq W\}$ to denote an $H \times W$ input image with $p_{i,j}$ being its pixel. Besides, we use

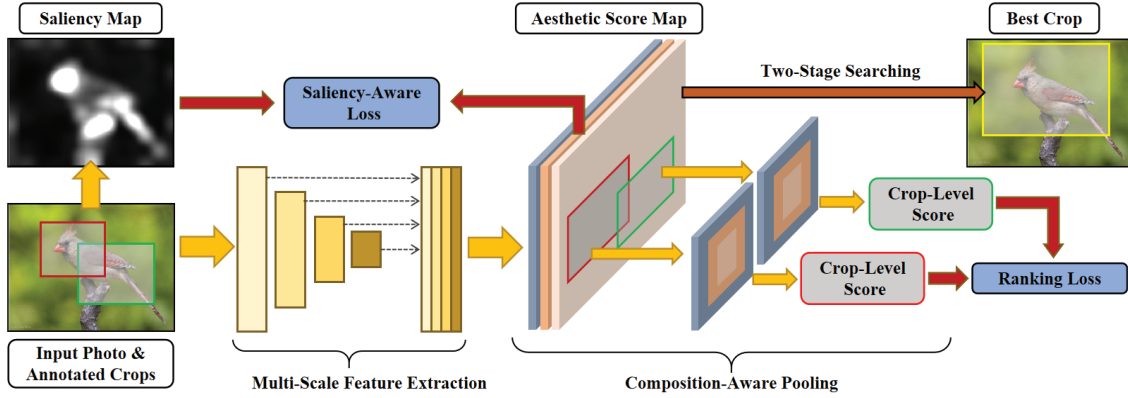


Figure 2: Illustration of our model. **Train stage:** We first obtain the aesthetic score map based on multi-scale features and then produce crop-level scores with composition-aware pooling. Finally, we train our model with ranking loss and saliency-aware loss. **Inference stage:** For each test image, we use the learned model to obtain its aesthetic score map. Then we apply two-stage searching to find the best image crop.

\mathbf{M} to represent the aesthetic score map of \mathcal{I} . At first, we assume that \mathbf{M} is a $H \times W$ matrix with $m_{i,j}$ being its element. Given an image crop $\mathcal{X}_k \subseteq \mathcal{I}$, we can obtain the aesthetic score of this crop via average pooling $\frac{1}{|\mathcal{X}_k|} \sum_{p_{i,j} \in \mathcal{X}_k} m_{i,j}$. Similarly, the aesthetic score of a region in this crop can also be calculated via average pooling. In this case, one region will get the same aesthetic score in different crops containing this region. However, as discussed in Section 1, we find that this design is not compatible with aesthetic evaluation. When evaluating the aesthetic quality of a region, we should consider its relative position in different crops.

Therefore, we propose composition patterns to achieve this goal. Several instantiations of composition patterns are listed in Figure 3, which share similar spirit with the auxiliary lines used in human-defined composition rules (see Figure 1). We assume that a composition pattern has L non-overlapping composition partitions indexed by $1, 2, \dots, L$. Then, given a crop and the composition pattern, we can divide this crop into L corresponding composition partitions based on the composition pattern. For each region in this crop, we can easily get its partition index (i.e., $1, 2, \dots, L$) based on its relative position in this crop. The aesthetic score of a region should be aware of its composition partition index in different crops, and we refer to this property as ‘‘composition-aware’’.

Inspired by R-FCN (Dai et al. 2016), we expand the aesthetic score map to a tensor $\mathbf{M} \in \mathcal{R}^{H \times W \times L}$ with one channel corresponding to one composition partition. Given an image crop $\mathcal{X}_k \subseteq \mathcal{I}$ and a pixel $p_{i,j} \in \mathcal{X}_k$, we use $\gamma_k(i, j) \in \{1, 2, \dots, L\}$ to denote the composition partition index of $p_{i,j}$ in \mathcal{X}_k . So the composition-aware score of $p_{i,j}$ in \mathcal{X}_k is $\phi_k(i, j) = m_{i,j,\gamma_k(i,j)}$.

Composition-Aware Pooling: To calculate crop-level aesthetic scores, we design a composition-aware pooling operation to realize this function. We use $\Phi(\cdot)$ to denote composition-aware pooling and thus $\Phi(\mathcal{X}_k)$ is the crop-level aesthetic score of \mathcal{X}_k . In particular, the crop-level aesthetic score $\Phi(\mathcal{X}_k)$ is defined as the average of the composition-




Figure 3: Four different composition patterns. **3x3 Grid:** 9 partitions of the same size. **Multi-rectangle:** 3 partitions, *center*, *middle*, and *outer*. As the white dotted lines show, it can be divided into a 6×6 grid in implementation. **Multi-direction:** 8 partitions for 8 directions. **Cross-rectangle:** 12 partitions, inspired by the camera auxiliary lines.

aware aesthetic scores of all pixels in the crop \mathcal{X}_k (see the visualization results in Figure 4):

$$\Phi(\mathcal{X}_k) \equiv \frac{1}{|\mathcal{X}_k|} \sum_{p_{i,j} \in \mathcal{X}_k} \phi_k(i, j) \quad (1)$$

In our implementation, we calculate (1) in an efficient way. Given a $h \times w$ crop \mathcal{X}_k , we first obtain its crop-level score map $\mathbf{M}_k \in \mathcal{R}^{h \times w \times L}$ from \mathbf{M} .

Suppose our used composition pattern is $x \times y$ grid pattern, we perform spatial average pooling within each cell of \mathbf{M}_k to get the pooled score map $\mathbf{M}'_k \in \mathcal{R}^{x \times y \times L}$. Note that $L = x \cdot y$ for $x \times y$ grid pattern, so \mathbf{M}'_k can be further reshaped into a square matrix $\mathbf{M}''_k \in \mathcal{R}^{L \times L}$. Then the crop-level score of \mathcal{X}_k is:

$$\Phi(\mathcal{X}_k) = \frac{1}{L} \sum \text{diag}(\mathbf{M}''_k), \quad (2)$$




Figure 4: The crop-level aesthetic score maps for three crops in the same image. The annotated scores are in black and the predictions are in red.

in which $\text{diag}(\cdot)$ denotes the diagonal elements.

When dealing with a non-grid pattern with L partitions, we first transform this pattern to a grid pattern with \tilde{L} partitions. In particular, we divide each partition into small cells and each small cell is a new partition (see the multi-rectangle in Figure 3 with $\tilde{L} = 36$). The multi-direction pattern is a little special but can also be handled in the same way. To be exact, we divide the whole pattern into 12×12 cells ($\tilde{L} = 144$) so that each partition is approximately covered by 18 cells. Then we use $\mathbf{M} \in \mathcal{R}^{H \times W \times L}$ to build a new map $\tilde{\mathbf{M}} \in \mathcal{R}^{H \times W \times \tilde{L}}$. If the \tilde{l} -th new partition comes from the l -th old partition, we have $m_{i,j,\tilde{l}} = m_{i,j,l}$, $\forall i, j$. Then we

directly use $\tilde{\mathbf{M}}$ to replace \mathbf{M} for composition-aware pooling. **Training with Ranking Loss:** We use $\{(\mathcal{X}_k, y_k) | \mathcal{X}_k \in \mathcal{I}, k = 1, \dots, K\}$ to denote the K annotated crops from \mathcal{I} , in which y_k is the annotated score of \mathcal{X}_k . During training, we need to learn $\Phi(\mathcal{X}_k)$ with supervision from y_k . Following (Chen et al. 2017b; Wei et al. 2018), we learn the relative ranking order of all the crops with the ranking loss, because the ability to correctly ranking crops with similar contents is essential in image cropping (Chen et al. 2017a).

For each training image, we first use its annotated crops to generate training pairs based on the differences of the annotated scores. Specifically, given a threshold δ , the crop pairs are denoted as

$$\mathcal{R} = \{(k, t) | y_k - y_t \geq \delta; k \neq t; \mathcal{X}_k, \mathcal{X}_t \in \mathcal{I}\}. \quad (3)$$

Since we have totally K annotated crops, we can generate at most $(K^2 - K)/2$ crop pairs with $\delta = 0$. By increasing δ , we can filter out confusing pairs that might be harmful to training. Then the remaining crop pairs are used to train our model with ranking loss:

$$\mathcal{L}_{\text{rank}} = \sum_{(k,t) \in \mathcal{R}} \max\{0, 1 + \Phi(\mathcal{X}_t) - \Phi(\mathcal{X}_k)\}. \quad (4)$$

3.2 Saliency-Aware Aesthetic Score Map

Before introducing visual saliency to our model, we first describe the concept of aesthetically important regions, which is helpful in interpreting image cropping.

Aesthetically Important Region: We use the term ‘‘aesthetically important’’ to denote the regions that are very influential in image cropping. We summarize two types of aesthetically important regions: 1) A region has important content which needs to be preserved, and thus a good crop is

supposed to include this region. These ‘‘content-important’’ regions will seriously influence the location of good crops. 2) A region is sensitive to its composition partition index. When including such ‘‘composition-sensitive’’ region in a crop, they should be carefully placed in proper positions; otherwise they may lead to aesthetic catastrophe.

In fact, content-important regions and composition-sensitive regions can be explained based on our aesthetic score map. We define two statistics $\theta_{\text{avg}}(\cdot)$ and $\theta_{\text{std}}(\cdot)$, in which $\theta_{\text{avg}}(p_{i,j})$ is the average aesthetic score of $p_{i,j}$ over all composition partitions and $\theta_{\text{std}}(p_{i,j})$ is the corresponding standard deviation:

$$\theta_{\text{avg}}(p_{i,j}) = \frac{1}{L} \sum_{l=1}^L m_{i,j,l}, \quad (5)$$

$$\theta_{\text{std}}(p_{i,j}) = \left(\frac{1}{L} \sum_{l=1}^L (m_{i,j,l} - \theta_{\text{avg}}(p_{i,j}))^2 \right)^{1/2}. \quad (6)$$

Intuitively, a region with high θ_{avg} has a large aesthetic score on average, and good crops are most likely to contain this region. So the regions with high θ_{avg} could be deemed as content-important regions. A region with θ_{std} is more sensitive to its composition partition index, because $\theta_{\text{std}}(\cdot)$ describes the average score difference caused by the change of its belonging composition partition. So the regions with high θ_{std} could be treated as composition-sensitive regions. After obtaining an aesthetic score map, we can draw heat maps based on θ_{avg} and θ_{std} to localize aesthetically important regions.

Saliency-Aware Loss: Visual saliency is a perceptual quality that can make objects attract more human attention (Hou and Zhang 2009). It is studied in different applications like saliency prediction (Pan et al. 2016), salient object detection (Borji et al. 2015), and video object segmentation (Wang et al. 2017). In the image cropping task, we can detect the salient regions to preserve the important contents in the best crop. Some previous methods (Wang and Shen 2017; Lu et al. 2019b) have already tried in this direction with a simple assumption: the most salient region should be a content-important region. On the premise of this assumption, their methods generate candidate crops covering the detected salient region, so that the predicted best crop will not miss the salient region. If our model follows this assumption, we should add a constraint that the salient regions have large θ_{avg} during training.

However, we realize that such an assumption has several limitations. First, many real-world images like a party photo could have multiple salient regions, and we cannot only focus on the most salient one. Second, a landscape photo may not have any salient object, which could result in failure cases as reported in (Lu et al. 2019b). Third, a salient object could also be a distraction that should be excluded from a good crop.

In our method, we adopt a more reasonable assumption: salient regions should be composition-sensitive, which can be explained as follows. According to the definition of saliency, a salient region must be visually attractive. When looking at an image, salient regions will draw more attention




Figure 5: Performance variation of our method with different trade-off parameter λ on the validation set of CPC dataset.

than others, so people will care more about whether they are placed in proper positions. We first use the saliency detection model (Hou and Zhang 2007) to obtain a saliency map \mathbf{S} for image \mathcal{I} , in which $s_{i,j} \in [0, 1]$ is the saliency score for $p_{i,j}$. Then, we propose our saliency-aware loss \mathcal{L}_{sal} by weighing $\theta_{\text{std}}(p_{i,j})$ with saliency score:

$$\mathcal{L}_{\text{sal}} = \frac{1}{HW} \sum_{p_{i,j} \in \mathcal{I}} (1 - s_{i,j}) \frac{\theta_{\text{std}}(p_{i,j})}{\bar{\theta}_{\text{std}}}, \quad (7)$$

in which $\bar{\theta}_{\text{std}}$ is the average of all $\theta_{\text{std}}(p_{i,j})$ in \mathcal{I} . For non-salient regions, $(1 - s_{i,j})$ is large, so they will be penalized for having large θ_{std} . For salient regions, $(1 - s_{i,j})$ is small, so they are allowed to have large θ_{std} . Notice that $\theta_{\text{std}}(\cdot)$ reflects our model capacity. In extreme case, if $\theta_{\text{std}}(p_{i,j}) = 0$ for all pixels, our aesthetic score tensor will reduce to aesthetic score matrix with only one channel, and the shrunk model capacity will cause the difficulty of minimizing $\mathcal{L}_{\text{rank}}$. By penalizing θ_{std} of non-salient regions, we actually encourage salient regions to have larger θ_{std} in order to fit the training data. In this way, salient objects are encouraged to be composition-sensitive, so that our model can learn to place salient objects in proper positions.

In the training stage, we use images with annotated crops and their saliency maps to train our ASM-Net with the total training loss:

$$\mathcal{L} = \mathcal{L}_{\text{rank}} + \lambda \mathcal{L}_{\text{sal}}, \quad (8)$$

in which λ is a trade-off parameter.

3.3 Two-Stage Search for Image Cropping

During inference, the saliency map is no longer needed. Given an image, we first use the learned model to produce an aesthetic score map. Based on the aesthetic score map, we aim to find the best crop by generating a lot of candidate crops and comparing their crop-level scores. Previous methods (Chen et al. 2017b; Wei et al. 2018; Zeng et al. 2019) adopt different strategies in this step. Here, we design a simple yet effective searching strategy to reduce the searching space. Specifically, our searching method has two stages. In the first stage, we conduct a coarse-grained search on a set of pre-defined anchor boxes as in (Wei et al. 2018). In the second stage, we conduct fine-grained searching around the top crops (*e.g.*, top 10) in the first stage, by slightly varying their center position, size, and aspect ratio.

Pattern	#Partitions	Accuracy \uparrow
1×1 grid	1	59.46
1×3 grid	3	75.64
3×1 grid	3	71.33
5×5 grid	25	86.38
multi-rectangle	5	81.34
multi-direction	8	83.22
cross-rectangle	12	87.59

Table 1: Accuracies (%) of our method with different composition patterns on the validation set of CPC dataset.

4 Experiments

4.1 Datasets

We use four image cropping datasets for experiments, two for training and two for testing.

CPC: The CPC dataset (Wei et al. 2018) is the first densely annotated image cropping dataset. It contains 10797 images, and each image has 24 crops with multiple scores annotated by Amazon Mechanical Turk workers. We use the average annotated score for each crop as ground truth.

GAICD: The GAICD dataset (Zeng et al. 2019) is also a densely annotated dataset, but each image has more annotated crops than the CPC dataset. It has 1236 images and about 86 crops for each image.

FCDB & FLMS: The FCDB (Chen et al. 2017a) and FLMS (Fang et al. 2014) datasets only have annotations for the best crop in each image, which are given by ten annotators. FCDB contains 358 images, and FLMS contains 500 images. These two datasets are only used as the test set in the image cropping task.

4.2 Implementation Details

Following the setting in (Wei et al. 2018; Zeng et al. 2019), we use VGG16 as the backbone model for feature extraction. For generating crop pairs, we have $\delta = 0.3$ for CPC dataset and $\delta = 0.9$ for GAICD dataset. We choose an unsupervised method in (Hou and Zhang 2007) to produce saliency maps for training images.

4.3 Ablation Study

We first conduct experiments to study the effectiveness of different composition patterns and our saliency-aware loss. Here, we train our model with the CPC dataset and leave 2000 images as the validation set. We report the ranking accuracy, *i.e.*, the accuracy of correctly ranking the generated crop pairs, on the validation set of CPC dataset.

Composition Pattern: We explore four types of composition patterns in Figure 3 and the results are reported in Table 1. First, we investigate the grid pattern with different numbers of partitions. When using 1×1 grid, our model is no longer composition-aware because it will assign a region with the same score regardless of its different composition partition indices in different crops. When using 1×3 (*resp.*, 3×1) grid, our model is only aware of horizontal (*resp.*, vertical) shift. We observe that 1×3 grid has higher accuracy, which means that horizontal shift has more impact on the

Method	IoU \uparrow	Disp \downarrow
RankSVM (Chen et al. 2017a)	0.6020	0.106
AesRankNet (Kong et al. 2016)	0.4843	0.140
MNA-CNN (Mai, Jin, and Liu 2016)	0.5042	0.136
VFN (Chen et al. 2017b)	0.6842	0.084
VPN (Wei et al. 2018)	0.7109	0.073
VEN (Wei et al. 2018)	0.7349	0.072
ASM-Net (ours)	0.7489	0.068

Table 2: Results of different methods on the FCDB dataset.

Method	IoU \uparrow	Disp \downarrow
Chen et al. (2016)	0.6400	0.075
Fang et al. (2014)	0.7400	-
Suh et al. (2003)	0.7200	0.063
ABP+AA (Wang and Shen 2017)	0.8100	0.057
VPN (Wei et al. 2018)	0.8352	0.044
VEN (Wei et al. 2018)	0.8365	0.041
ASM-Net (ours)	0.8486	0.039

Table 3: Results of different methods on the FLMS dataset.

aesthetic quality of a crop, which has also been proved by empirical studies (Palmer, Schloss, and Sammartino 2013; Abeln et al. 2016). Then, we compare four types of composition patterns and report the best results for four types of composition patterns. We conjecture that their performance differences are attributed to their sensitivity to different types of position shifts. To be exact, multi-rectangle and multi-direction are only sensitive to one type of position shift (*i.e.*, in-out shift or directional shift), which restrains them from capturing more complicated composition rules in the training data, so they have lower results. The best result comes from the cross-rectangle, which is like a combination of multi-rectangle and multi-direction. It is inspired by camera auxiliary lines, which proves to be more compatible with learning composition rules. Due to the competitive performance of the cross-rectangle, we use it as the default composition pattern for the rest of the experiments.

Saliency-Aware Loss: Recall that we have a trade-off parameter λ before our saliency-aware loss in (8). We vary λ in the range of $[0, 100]$ and report the results in Figure 5. When we do not use the saliency-aware loss ($\lambda=0$), the accuracy drop is 3.22% compared with the best accuracy achieved when $\lambda=0.1$. When λ becomes larger, the accuracy begins to drop because it will reduce our model capacity, as discussed in Section 3.2. We use $\lambda=0.1$ by default for the rest of experiments.

4.4 Performance on Image Cropping

Image cropping models are usually evaluated based on single best crop prediction. However, Zeng et al. argue that image cropping is naturally a subjective and flexible task without a unique solution, and thus suggest evaluation based on multiple crop ranking. Here we conduct experiments based on both best crop prediction and multiple crop ranking.

Predicting Best Crop: In this task, we predict the best crop and compare it with the ground truth crop. We follow the ex-

Method	SRCC \uparrow	Acc $_5$ \uparrow	Acc $_{10}$ \uparrow
VFN (Chen et al. 2017b)	0.450	26.7	38.7
VEN (Wei et al. 2018)	0.621	37.6	50.9
GAIC (Zeng et al. 2019)	0.735	50.2	68.5
ASM-Net (ours)	0.766	54.3	71.5

Table 4: Results of different methods on the GAICD dataset.

periment setting in (Chen et al. 2017b; Wei et al. 2018) and use two metrics for performance evaluation: intersection-over-union (IoU) and boundary displacement (Disp). We use the CPC dataset as the training set, and FCDB and FLMS datasets as two test sets. The results of our ASM-Net and baseline methods are reported in Table 2&3. On FLMS, we beat the saliency-based approach (Wang and Shen 2017). Our method achieves the best performance on both FCDB and FLMS datasets.

Ranking Multiple Crops: The above experiments are designed for locating a single best crop, where only the prediction of the best crop is evaluated. In the densely annotated dataset, each image has multiple annotated crops, so Zeng et al. (2019) proposed to compare the ranking order of all annotated crops with two different metrics: Spearman’s rank-order correlation coefficient (SRCC) and generalized top-N accuracy (Acc_N). While IoU and Disp only care about the best crop prediction, SRCC and Acc_N evaluate the performance of correctly ranking multiple annotated crops. Following (Zeng et al. 2019), we choose 1036 images in GAICD dataset for training and the remaining 200 images for testing. Many previous methods can only predict the best crop, so they are not applicable to this task. The results are summarized in Table 4, from which we observe that our method beats all the applicable baseline methods.

Discussion: The main advantage of our method is that the aesthetic score map is both composition-aware and saliency-aware. Remember that a good image crop is achieved by placing important content with a good composition, so these two properties are very important in the image cropping task.

With saliency-aware property, our model can focus on the salient regions, which are more likely to contain important content. Compared with the saliency-based approach ((Wang and Shen 2017)), our method has a more reasonable assumption that is more compatible with aesthetic evaluation. With composition-aware property, our model is able to encode complicated composition rules learned from training data. Especially when ranking the overlapping crops which are visually similar, our method gains an edge by referring to the composition partition index of their common content. In summary, our model is more compatible with the image cropping task, so it can make better use of the training data and have better performance on both tasks.

4.5 Application

Localizing Aesthetically Important Region: Following the discussion in Section 3.2, we locate aesthetically important regions by drawing heat maps with $\theta_{\text{avg}}(\cdot)$ and $\theta_{\text{std}}(\cdot)$ in Figure 6. We compare the heat maps between $\lambda=0$ and $\lambda=0.1$. We can observe that when not using saliency-aware

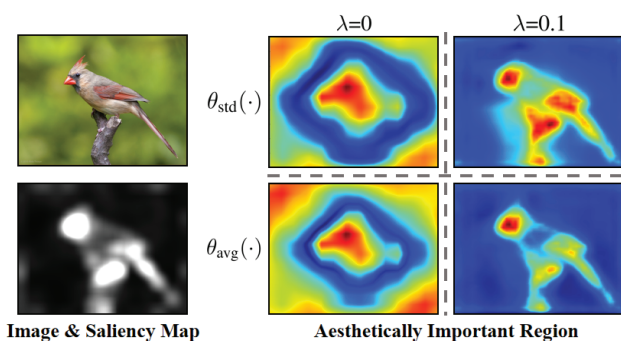


Figure 6: Illustration of aesthetically important regions based on the heat maps of $\theta_{std}(\cdot)$ and $\theta_{avg}(\cdot)$ with $\lambda=0$ and $\lambda=0.1$.




Figure 7: Illustration of spatial bias. High-score crops are prone to cover four corners of the whole image.

loss ($\lambda=0$), the image corners are aesthetically important. We believe it is caused by the spatial bias in the training images.

The spatial bias means that the annotated score of a crop is related to its position in the whole image. For example, since most photos will place the important contents at the center, the high-score crops are usually central and large enough to cover them, so they are more likely to cover the four corners of the image compared with low-score crops (see Figure 7). So the model is prone to assign high aesthetic scores to the corners of the image. In fact, a previous method GAIC has leveraged such statistical prior information to reduce low-score crops in candidate crop generation. By using saliency-aware loss, we are actually preventing the model from overfitting such spatial bias, so the aesthetically important regions are closer to the salient regions.

Diverse Image Cropping: To demonstrate the ability of our model, we visualize our cropping results in three typical scenes: one salient object, multiple salient objects, and no salient object. We first show the best cropping result in three typical scenes (two images for each scene) with five aspect ratios (9:16, 3:4, 1:1, 4:3, and 16:9) in Figure 8 (see more results in Supplementary). Notice that a saliency-based method might fail in the scene with no salient object, as reported in (Lu et al. 2019b). However, the results have demonstrated that our model can generally produce satisfactory results in different scenes with different aspect ratios.

Circular Cropping: Since we obtain crop-level aesthetic score by pooling the aesthetic score map, our model can be naturally used to evaluate image crops of arbitrary shape, by slightly changing the composition-aware pooling operation.

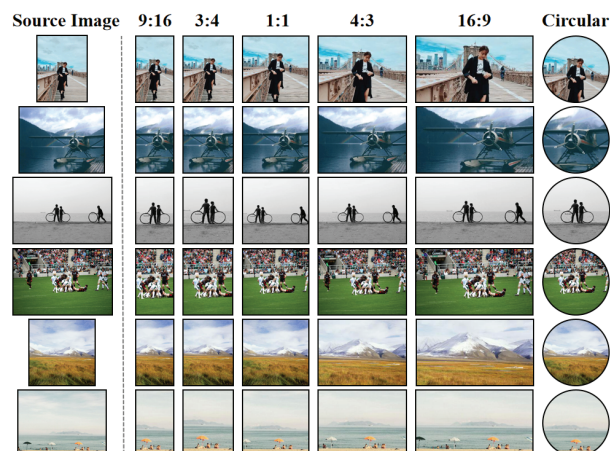


Figure 8: Image cropping results on six images (left column) with different aspect ratios and shapes.

Now we extend the rectangle cropping task to circular cropping since it can be used in real-world applications like generating circular user icons for websites. We use $\mathcal{X}_t \subset \mathcal{I}$ to denote a candidate circular crop, which is also the inscribed circle of a square crop \mathcal{X}_k . During training, we still use rectangle crops due to the lack of annotated circular crops. During prediction, we use $\phi_k(i, j)$ from \mathcal{X}_k as the composition-aware score of $p_{i, j}$ but only pool from the inscribed circle \mathcal{X}_t , so the circular crop-level score of \mathcal{X}_t is :

$$\Phi'(\mathcal{X}_t) = \frac{1}{|\mathcal{X}_t|} \sum_{p_{i, j} \in \mathcal{X}_t} \phi_k(i, j). \quad (9)$$

With the circular crop-level score defined above, our model finds some best circular crops, which are shown in Figure 8 (see more results in Supplementary). The results demonstrate the generality of our proposed model. Since we have not trained the model with any annotated circular image crops, it also proves that our model has mastered a universal aesthetic evaluation criterion that can be shared among crops with different shapes.

5 Conclusions

In this paper, we have proposed an interpretable model to unveil the intrinsic mechanism of the image cropping model. Our model can produce an aesthetic score map, in which the aesthetic score of a region is aware of its composition partition index in different crops. Moreover, by introducing visual saliency, our model can learn to place salient objects in proper positions. Our model has achieved the best performance in image cropping tasks and demonstrated good generality in real-world applications.

Acknowledgement

The work is supported by the National Key R&D Program of China (2018AAA0100704) and is partially sponsored by Shanghai Sailing Program (BI0300271) and Startup Fund for Youngman Research at SJTU (WF220403041).

References

- Abeln, J.; Fesz, L.; Amirshahi, S. A.; McManus, I. C.; Koch, M.; Kreysa, H.; and Redies, C. 2016. Preference for well-balanced saliency in details cropped from photographs. *Frontiers in human neuroscience* 9:704.
- Borji, A.; Cheng, M.-M.; Jiang, H.; and Li, J. 2015. Salient object detection: A benchmark. *IEEE transactions on image processing* 24(12):5706–5722.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2015. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*.
- Chen, J.; Bai, G.; Liang, S.; and Li, Z. 2016. Automatic image cropping: A computational complexity study. In *CVPR*, 507–515.
- Chen, Y.-L.; Huang, T.-W.; Chang, K.-H.; Tsai, Y.-C.; Chen, H.-T.; and Chen, B.-Y. 2017a. Quantitative analysis of automatic image cropping algorithms: A dataset and comparative study. In *WACV*, 226–234. IEEE.
- Chen, Y.-L.; Klopp, J.; Sun, M.; Chien, S.-Y.; and Ma, K.-L. 2017b. Learning to compose with professional photographs on the web. In *ACM-Multimedia*, 37–45. ACM.
- Dai, J.; Li, Y.; He, K.; and Sun, J. 2016. R-fcn: Object detection via region-based fully convolutional networks. In *NIPS*, 379–387.
- Esmaeili, S. A.; Singh, B.; and Davis, L. S. 2017. Fast-at: Fast automatic thumbnail generation using deep neural networks. In *CVPR*, 4622–4630.
- Fang, C.; Lin, Z.; Mech, R.; and Shen, X. 2014. Automatic image cropping using visual composition, boundary simplicity and content preservation models. In *ACM-Multimedia*, 1105–1108. ACM.
- Hosu, V.; Goldlucke, B.; and Saupe, D. 2019. Effective aesthetics prediction with multi-level spatially pooled features. In *CVPR*, 9375–9383.
- Hou, X., and Zhang, L. 2007. Saliency detection: A spectral residual approach. In *CVPR*, 1–8. Ieee.
- Hou, X., and Zhang, L. 2009. Dynamic visual attention: Searching for coding length increments. In *NIPS*, 681–688.
- Jiang, M.; Huang, S.; Duan, J.; and Zhao, Q. 2015. Salicon: Saliency in context. In *CVPR*, 1072–1080.
- Kong, S.; Shen, X.; Lin, Z.; Mech, R.; and Fowlkes, C. 2016. Photo aesthetics ranking network with attributes and content adaptation. In *ECCV*, 662–679. Springer.
- Li, D.; Wu, H.; Zhang, J.; and Huang, K. 2018. A2-rl: Aesthetics aware reinforcement learning for image cropping. In *CVPR*, 8193–8201.
- Li, X.; Li, X.; Zhang, G.; and Zhang, X. 2019. Image aesthetic assessment using a saliency symbiosis network. *Journal of Electronic Imaging* 28(2):023008.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *ECCV*, 21–37. Springer.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*, 3431–3440.
- Lu, X.; Lin, Z.; Jin, H.; Yang, J.; and Wang, J. Z. 2014. Rapid: Rating pictorial aesthetics using deep learning. In *ACM-Multimedia*, 457–466. ACM.
- Lu, X.; Lin, Z.; Shen, X.; Mech, R.; and Wang, J. Z. 2015. Deep multi-patch aggregation network for image style, aesthetics, and quality estimation. In *ICCV*, 990–998.
- Lu, P.; Zhang, H.; Peng, X.; and Jin, X. 2019a. An end-to-end neural network for image cropping by learning composition from aesthetic photos. *arXiv preprint arXiv:1907.01432*.
- Lu, P.; Zhang, H.; Peng, X.; and Peng, X. 2019b. Aesthetic guided deep regression network for image cropping. *Signal Processing: Image Communication*.
- Lu, W.; Xing, X.; Cai, B.; and Xu, X. 2019c. List-wise view ranking for image cropping. *arXiv preprint arXiv:1905.05352*.
- Mai, L.; Jin, H.; and Liu, F. 2016. Composition-preserving deep photo aesthetics assessment. In *CVPR*, 497–506.
- Murray, N.; Marchesotti, L.; and Perronnin, F. 2012. Ava: A large-scale database for aesthetic visual analysis. In *CVPR*, 2408–2415. IEEE.
- Palmer, S. E.; Schloss, K. B.; and Sammartino, J. 2013. Visual aesthetics and human preference. *Annual review of psychology* 64:77–107.
- Pan, J.; Sayrol, E.; Giro-i Nieto, X.; McGuinness, K.; and O’Connor, N. E. 2016. Shallow and deep convolutional networks for saliency prediction. In *CVPR*, 598–606.
- Suh, B.; Ling, H.; Bederson, B. B.; and Jacobs, D. W. 2003. Automatic thumbnail cropping and its effectiveness. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, 95–104. ACM.
- Sun, J., and Ling, H. 2013. Scale and object aware image thumbnailing. *International journal of computer vision* 104(2):135–153.
- Wang, W., and Shen, J. 2017. Deep cropping via attention box prediction and aesthetics assessment. In *ICCV*, 2186–2194.
- Wang, W.; Shen, J.; Yang, R.; and Porikli, F. 2017. Saliency-aware video object segmentation. *IEEE transactions on pattern analysis and machine intelligence* 40(1):20–33.
- Wei, Z.; Zhang, J.; Shen, X.; Lin, Z.; Mech, R.; Hoai, M.; and Samaras, D. 2018. Good view hunting: Learning photo composition from dense view pairs. In *CVPR*, 5437–5446.
- Yan, J.; Lin, S.; Bing Kang, S.; and Tang, X. 2013. Learning the change for automatic image cropping. In *CVPR*, 971–978.
- Zeng, H.; Li, L.; Cao, Z.; and Zhang, L. 2019. Reliable and efficient image cropping: A grid anchor based approach. In *CVPR*, 5949–5957.
- Zhang, L.; Song, M.; Yang, Y.; Zhao, Q.; Zhao, C.; and Sebe, N. 2013. Weakly supervised photo cropping. *IEEE Transactions on Multimedia* 16(1):94–107.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *CVPR*, 2921–2929.