

Adaptive Fractional Dilated Convolution Network for Image Aesthetics Assessment

Qiuyu Chen¹, Wei Zhang², Ning Zhou³, Peng Lei³, Yi Xu², Yu Zheng⁴, Jianping Fan¹

¹Department of Computer Science, UNC Charlotte

²Shanghai Key Laboratory of Intelligent Information Processing, School of Computer Science, Fudan University

³Amazon Lab126

⁴School of Cyber Engineering, Xidian University

{qchen12, jfan}@uncc.edu, {weizh, yxu17}@fudan.edu.cn

{ningzho, leipeng}@amazon.com, yuzheng.xidian@gmail.com

Abstract

To leverage deep learning for image aesthetics assessment, one critical but unsolved issue is how to seamlessly incorporate the information of image aspect ratios to learn more robust models. In this paper, an adaptive fractional dilated convolution (AFDC), which is aspect-ratio-embedded, composition-preserving and parameter-free, is developed to tackle this issue natively in convolutional kernel level. Specifically, the fractional dilated kernel is adaptively constructed according to the image aspect ratios, where the interpolation of nearest two integer dilated kernels are used to cope with the misalignment of fractional sampling. Moreover, we provide a concise formulation for mini-batch training and utilize a grouping strategy to reduce computational overhead. As a result, it can be easily implemented by common deep learning libraries and plugged into popular CNN architectures in a computation-efficient manner. Our experimental results demonstrate that our proposed method achieves state-of-the-art performance on image aesthetics assessment over the AVA dataset [18].

1. Introduction

This paper addresses image aesthetics assessment where the goal is to predict the given image an aesthetic score. Automatic image aesthetics assessment has many applications such as album photo recommendation, auxiliary photo editing, and multi-shot photo selection. The task is challenging because it entails computations of both global cues (e.g. scene, exposure control, color combination, etc) and localization information (composition, photographic angle, etc).

Early approaches extract aesthetic features according to photographic rules (lighting, contrast) and global image composition (symmetry, rule of thirds), which require ex-

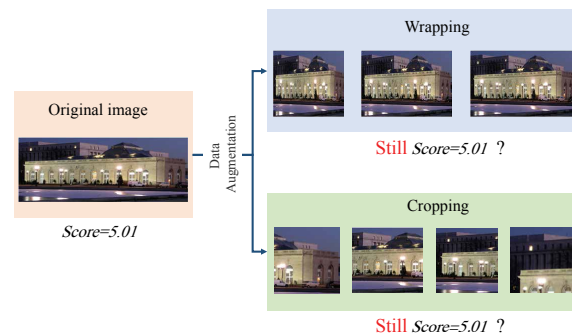


Figure 1: Image warping and cropping are widely used for data augmentation, but they alter the object aspect ratios and composition, causing different aesthetics perceptions. Assigning the groundtruth aesthetic score of the original image to the altered image may introduce label noise and deteriorate the discriminative ability.

tensive manual designs [3, 5, 13, 19, 25, 28]. However, manual design for such aesthetic features is not a trivial task even for experienced photographers. Recent work adopts deep convolutional neural networks for image aesthetics assessment by learning models in an end-to-end fashion. The models mainly use three types of formulations: binary classification labels [12, 15, 16, 29, 20, 33, 23], scores [17, 27, 8], and rankings [14, 22].

In the aforementioned methods, the backbone networks are usually adopted from an image classification network. The data augmentation methods, *i.e.* image cropping and warping, are widely used for preventing overfitting in the image recognition task. However, a shortcoming is that the compositions and object aspect ratios are altered, which may introduce label noise and harm the task of aesthetics assessment (Fig. 1). A succinct solution proposed in MNA-CNN [16] is to feed one original-size image into the network at a time during training and test (bottom stream in Fig. 2). A major constraint of the approach is that im-

ages with different aspect ratios cannot be concatenated into batches because the aspect ratio of each image should be preserved. Thus it slows down the training and inference.

In this paper, we aim to develop a novel adaptive fractional dilated convolution that is mini-batch compatible. As shown in the top row in Fig. 2, our network adaptively dilates the convolution kernels to the composition-preserving warped images according to the image aspect ratios such that the effective receptive field of each dilated convolution kernel is the same as the regular one. Specifically, as illustrated in Fig. 3, the fractional dilated convolution kernel is adaptively interpolated by the nearest two integer dilated kernels with the same kernel parameters. Thus no extra learning parameters are introduced.

The **benefits** of our method can be summarized as follows: (a) By embedding the information of aspect ratios to construct the convolution layers adaptively, it can explicitly relate the aesthetic perception to the image aspect ratios while preserving the composition; (b) It is parameter-free and thus can be easily plugged into the popular network architectures; (c) Through the deduction, we show that our proposed method can be mini-batch compatible and easily implemented by common deep learning libraries (*e.g.* PyTorch, Tensorflow); (d) A grouping strategy is introduced to reduce the computational overhead for efficient training/inference; (e) We achieve state-of-the-art performance for image aesthetics assessment on the AVA dataset [18].

2. Related Work

In this section, we provide a brief review of some of the most relevant works on: (a) image aesthetics assessment; (b) preserving image aspect ratios and compositions; (c) dilated convolution; (d) dynamic kernels.

Image Aesthetics Assessment. The existing methods on image aesthetics assessment can be mainly categorized into three formulations: (1) *Binary (or mean) aesthetic label*: Kao *et al.* [12] propose a multi-task CNN, A&C CNN, which jointly learns both the category classification and the aesthetic perception. Mai *et al.* [16] address the composition problem in image aesthetics assessment and aggregates multiple sub-networks with different sizes of adaptive pooling layer. Ma *et al.* [15] feed the patches sampled from the saliency map of the original image into VGG16 [24] with an aggregation layer, where a layer-aware subnet considering path localizations is leveraged to get the final prediction. Sheng *et al.* [23] assign adaptively larger weights to meaningful training cropping patches according to the prediction errors during the training and aggregate the multi-patch predictions during the test. Hosu *et al.* [8] propose to incorporate multi-level spatially pooled features from the intermediate blocks in a computation efficient manner. (2) *Ranking score*: Instead of classification or regression formulations, a joint loss of Euclidean and ranking [14] is proposed and

a triplet ranking loss [22] is developed. (3) *Score distribution*: To address the ordered score distribution, Hossein Talebi and Peyman Milanfar [27] introduce Earth Mover’s Distance as a loss function to train 10-scale score distribution. Since the image aesthetics is a subjective property and outlier opinions may appear, Naila Murray and Albert Gordo [17] introduce Huber Loss to train 10-scale score distribution. Besides using the mean score of multiple raters, Ren *et al.* [20] propose a sub-network to learn a personal rating offset along with the generic aesthetic network and output the personalized score prediction.

Preserving Image Aspect Ratios and Compositions. Multi-patch sampling over the original images is used to preserve the aspect ratios and proves to be effective [15, 23, 8]. A major concern is that sampling patches from the original image may alter essential aesthetic factors (color histogram, object-background ratio) of the original image and the complete aesthetics features are lost. In contrast, our proposed method adaptively restores the original receptive fields from the composition-preserving warping images in an end-to-end fashion. The approach of MNA-CNN [16] is the most related to ours, as they proposed to preserve image aspect ratios and compositions by feeding the original image into the network, one at a time. A major constraint of the approach is that images with different aspect ratios cannot be concatenated into batches because the aspect ratio of each image should be preserved. Thus it tends to slow down the training and inference processes. On the other hand, our proposed method is mini-batch compatible and can be easily implemented by common deep learning libraries.

Dilated Convolution. Our adaptive fractional dilated convolution is motivated by the dilated convolution [31] and atrous convolution [1] in semantic segmentation, but it differs from them in several aspects: (1) Our adaptive fractional dilated convolution is to restore the receptive fields for warped images to the same as regular convolution for original images, while dilated convolution is proposed to retrain the large receptive without down-sampling. (2) The dilation rate can be fractional in our method. (3) The construction of fractional dilated kernel is dynamic respecting the aspect ratios.

Dynamic Kernels. Deformable convolution [2] is proposed to construct the receptive fields dynamically and adaptively by learning better sampling in the convolutional layer. Our proposed method differs from deformable convolution in two folds: (a) The deformable convolution is proposed to learn better sampling in the convolutional layers, whereas our method adapts the receptive fields into the original aspect ratios. Therefore, our proposed method is parameter-free while the deformable convolution requires parameterized layers to predict the sampling indices. (b) Our method provides a concise formula for mini-batch training and it can be easily implemented by the common deep learning

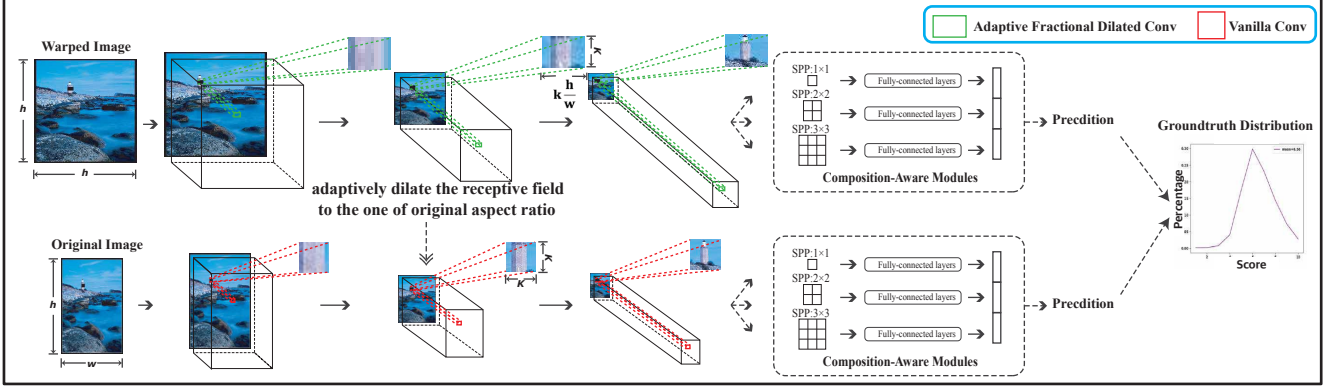


Figure 2: **Overview of adaptive fractional dilated CNN (above) and the comparison with vanilla CNN (below): Each fractional dilated Conv (above) operated on wrapped input adaptively dilates the same receptive field as the vanilla Conv (below) operated on the original image. It thus helps with the problems: (a) Becomes mini-batch compatible by composition-preserving warping instead of feeding original-size image (b) Preserves aesthetic features related to aspect ratios by adaptive kernel dilation.**

frameworks. On the other hand, the deformable convolution needs to rewrite the convolution operation in CUDA and tends to be slow due to the indexing operation.

3. Adaptive Fractional Dilated Convolution

In this section, we first introduce the adaptive kernel interpolation to tackle the misalignment due to fractional sampling in the proposed method. We then derive a concise formulation for it in the setting of mini-batch and discuss their computational overhead. Finally, we describe the loss function and an additional composition-aware structure for the composition-preserving warping batch.

3.1. Adaptive Kernel Interpolation

As stated in Section 1, cropping modifies the composition of the original image and causes the loss of some critical aesthetics information. As a result, image cropping introduces somewhat label noises in the training stage. To preserve the composition, we firstly warp the image into a fixed size. For network training, such a simple image warping approach suffers from the problem of overfitting due to the absence of data augmentation. Motivated by SPP [6], we adopt random-size warping during the training stage and feed the mini-batch into the networks with global pooling or SPP modules, which can naturally handle arbitrary-size batch inputs. Overall, the random-size warping provides effective data augmentation for training scale-invariant networks while preserving the image compositions.

To cope with the distortion induced by warping, the receptive field of the convolution kernel should be consistent with the receptive field of the convolution kernel that is operated on the image with original aspect ratio. Our proposed approach tackles the distortion issue by adaptively dilating the kernels to the original aspect ratio, as illustrated

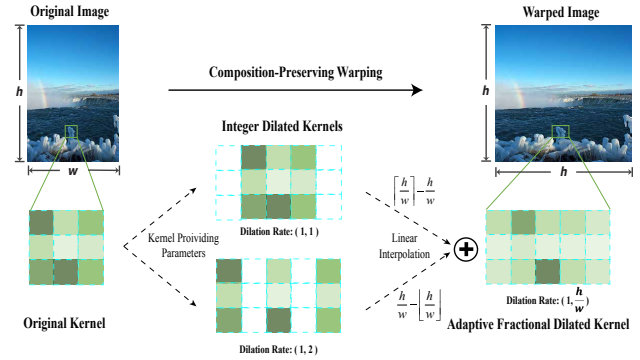


Figure 3: **Illustration of kernel interpolation: linear interpolation of the nearest two integer dilated kernels shared same kernel parameters are used to tackle the sampling misalignment from fractional dilation rates.**

in Fig. 2. Since the aspect ratio could be fractional, the dilation rate could be a fraction as well. To tackle the misalignment of feature sampling, we use the linear interpolation of two nearest integer dilation rates to construct the fractional dilation kernel.

Suppose that w and h represent the width and height of original images, respectively. If $h > w$ and $\frac{h}{w}$ is not a integer, as illustrated in Fig. 3, AFDC (adaptive fractional dilated convolution) kernel k_{AFDC}^n in n -th layer is constructed as:

$$k_{AFDC}^n = (\lceil r \rceil - r)k_{(1, \lceil r \rceil)}^n + (r - \lfloor r \rfloor)k_{(1, \lfloor r \rfloor)}^n \quad (1)$$

where $r = \frac{h}{w}$. For any non-integer r , it is in the interval $[\lfloor r \rfloor, \lceil r \rceil]$ whose length is equal to 1. $\lfloor r \rfloor$ and $\lceil r \rceil$ are two integers nearest to r . $k_{(1, \lfloor r \rfloor)}^n$ and $k_{(1, \lceil r \rceil)}^n$ are two dilated kernels with the nearest integer dilation rates $\lfloor r \rfloor$ and $\lceil r \rceil$ for n th layer, respectively. More specifically, as shown in Fig. 3, $r \in [1, 2]$, $\lfloor r \rfloor = 1$, $\lceil r \rceil = 2$. We note that both

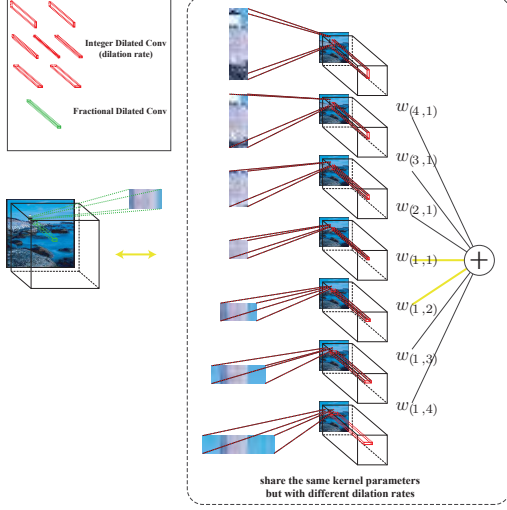


Figure 4: **Illustration for mini-batch compatibility: the distributive property of convolution operation (c.f. Eq. (3)) makes the fractional dilated conv easily implemented and compatible for mini-batch computation with a zero-padded weight vector/matrix (c.f. Eq. (5))**

$k_{(1,1)}^n$ and $k_{(1,2)}^n$ inherit the same learning parameters from the original kernel.

Likewise, if $w > h$ and $\frac{w}{h}$ is not an integer, then we choose:

$$k_{AFDC}^n = (\lceil r \rceil - r)k_{(\lceil r \rceil, 1)}^n + (r - \lfloor r \rfloor)k_{(\lfloor r \rfloor, 1)}^n \quad (2)$$

If $r = \frac{h}{w}$ is an integer, it is enough for us to employ integer dilated kernel.

Therefore, the fractional dilated kernel is adaptively constructed for each image with respect to w and h as shown in Fig. 3. In addition, all the integer dilation kernels share the same kernel parameters and thus no extra learning parameters are introduced.

3.2. Mini-Batch Computation and Implementation

To implement the dynamic kernel interpolation in Eq. (1) and Eq. (2) directly, we need to rewrite the kernel-level code due to the diverse kernels in mini-batch. However, through the following deduction, we show that the proposed method can be easily implemented by common deep learning libraries, e.g. PyTorch and TensorFlow.

Using the distributive property of convolution operation, the transformation of the feature maps generated by the adaptive fractional dilated Conv kernels in Eq. (1) can be formulated as:

$$\begin{aligned} f_{n+1} &= k_{AFDC}^n * f_n \\ &= \left[(\lceil \frac{w}{h} \rceil - \frac{w}{h})k_{(1, \lfloor \frac{w}{h} \rfloor)}^n + (\frac{w}{h} - \lfloor \frac{w}{h} \rfloor)k_{(1, \lceil \frac{w}{h} \rceil)}^n \right] * f_n \quad (3) \\ &= (\lceil \frac{w}{h} \rceil - \frac{w}{h})k_{(1, \lfloor \frac{w}{h} \rfloor)}^n * f_n + (\frac{w}{h} - \lfloor \frac{w}{h} \rfloor)k_{(1, \lceil \frac{w}{h} \rceil)}^n * f_n \end{aligned}$$

where f_n denotes the feature maps for the n th layer and $*$ denotes convolution.

In mini-batch training and inference, we can construct multiple kernels with different dilation rates ($rate_k^i, rate_k^j$) from the same kernel parameters and then use a zero-padded interpolation weight vector \mathbf{w} to compute the operation adaptively for each image as:

$$\begin{aligned} f_{n+1} &= k_{AFDC}^n * f_n \\ &= \sum_k w_{(rate_k^i, rate_k^j)} k_{(rate_k^i, rate_k^j)}^n * f_n \quad (4) \\ &= \tilde{\mathbf{w}} \mathbf{f}_n \end{aligned}$$

which is just the inner product of two vectors:

$$\mathbf{w} = [w_{(rate_1^i, rate_1^j)}, \dots, w_{(rate_K^i, rate_K^j)}] \quad (5)$$

and

$$\tilde{\mathbf{f}}_n = [k_{(rate_1^i, rate_1^j)}^n * f_n, \dots, k_{(rate_K^i, rate_K^j)}^n * f_n]^\top \quad (6)$$

where the number of dilation kernels is K . As shown in Fig. 4, the interpolation weight $w_{(rate_k^i, rate_k^j)}$ for each instance is either $w_{(rate_k^i, 1)}$ or $w_{(1, rate_k^j)}$, defined as follows:

$$\begin{aligned} w_{(rate^i, 1)} &= \begin{cases} r - (rate^i - 1), & \text{if } rate^i - r \in [0, 1) \\ (rate^i + 1) - r, & \text{if } rate^i - r \in (-1, 0) \\ 0, & \text{else} \end{cases} \\ w_{(1, rate^j)} &= \begin{cases} r - (rate^j - 1), & \text{if } rate^j - r \in [0, 1) \\ (rate^j + 1) - r, & \text{if } rate^j - r \in (-1, 0) \\ 0, & \text{else} \end{cases} \end{aligned} \quad (7)$$

In mini-batch, suppose that batch size is B , then the $n+1$ th feature maps \mathbf{F}_{n+1} can be formulated as:

$$\mathbf{F}_{n+1} = [\mathbf{f}_{n+1}^1, \dots, \mathbf{f}_{n+1}^B] = [\mathbf{w}^1 \tilde{\mathbf{f}}_n^1, \dots, \mathbf{w}^B \tilde{\mathbf{f}}_n^B] \quad (8)$$

The computation of the above $[\tilde{\mathbf{f}}_n^1, \dots, \tilde{\mathbf{f}}_n^B]$ can be done efficiently in the mini-batch as:

$$[k_{(rate_1^i, rate_1^j)}^n * \mathbf{F}_n, k_{(rate_2^i, rate_2^j)}^n * \mathbf{F}_n, \dots, k_{(rate_K^i, rate_K^j)}^n * \mathbf{F}_n]^\top \quad (9)$$

We note that the activation function and batch normalization are omitted in the formulas for concise illustration.

The formula in Eq. (8) can be interpreted as a dot production followed by a sum reduction between interpolation weight matrix \mathbf{W} and Eq. (9), which thus can be efficiently implemented by common deep learning frameworks (Pytorch, Tensorflow, etc.). Each integer dilated Conv, $k_{(rate_k^i, rate_k^j)}^n * \mathbf{F}_n$ in Eq. (9), is computed as a normal dilated Conv layer with the shared learning parameters.



Figure 5: **Grouping strategy to reduce computational overhead: The integer dilated Convs can be shared by properly grouped images according to aspect ratios.**

Network	#Params	#Multi-Adds	Speed (train)	Speed (test)
VGG16	138M	15.3G	8.14 it/s	12.91 it/s
2-dilation	138M	30.7G	2.70 it/s	3.85 it/s
7-dilation	138M	109.1G	0.73 it/s	0.93 it/s
ResNet50	25.6M	3.5G	12.49 it/s	22.80 it/s
2-dilation	25.6M	5.6G	8.32 it/s	14.81 it/s
2-dilation*	25.6M	6.5G	6.20 it/s	9.88 it/s
7-dilation	25.6M	10.6G	3.22 it/s	5.28 it/s
7-dilation*	25.6M	18.8G	2.08 it/s	3.12 it/s

Table 1: **Computation comparison: training batch size is set to 16, test batch size is set to 32. The speed is the average result for 100 iterations from the test on single GTX 1080Ti. The fractional dilated Conv is embedded for all BottleNets in ResNet50 while * denotes additional embedding dilation for the first 7×7 Conv layer as well.**

Computational overhead The computational overhead is determined by the number of integer dilated kernels and the number of convolutional layers whose kernel sizes are not 1×1 . As shown in Table 1, the BottleNet in ResNet50 [7] contains two 1×1 kernels and one 3×3 kernel. Since only 3×3 kernel introduces the computational overhead, the computational cost for 2 integer dilations is roughly 1.5 times of the original model, while VGG16 [24] consists of the majority of 3×3 kernels and thus the computation cost is approximately 2 times. Some additional computational overhead is caused by the interpolation operation of different dilation kernels.

Reducing overhead with a grouping strategy In practice, the aspect ratios, $\frac{w}{h}$, of most of images would fall into $[\frac{1}{2}, 2]$, e.g. 97.8% of the training and testing images in the AVA [18] dataset. Training efficiency can be optimized by grouping batches, e.g. training with three dilation kernels for the most batches, $DilationRates = \{(2, 1), (1, 1), (1, 2)\}$ for the images whose aspect ratios fall into $[\frac{1}{2}, 2]$. For the datasets with more diverse aspect ratios, a more fine-grained grouping strategy could be applied. As illustrated in Fig. 5, images with aspect ratio range $[4, 3]$ (above) and $[\frac{1}{2}, 1]$ (below) share the valid integer dilated Convs in the grouped batches.

Parallel optimization The calculation of multiple integer dilated kernels in each convolutional layer is equivalent to

broadening the output channel size by the number of dilation kernels. In another words, the computation of dilated Conv group, $\{k_{(rate_k^i, rate_k^j)}^n * F_n\}$, can be optimized through parallel computing. WideResNet [32] claims that increasing the width of Conv layers is more accommodating to the nature of GPU computation and helps effectively balance computations more optimally. However, from Table 1, the actual training and testing speeds are approximately linearly correlated with # Multi-Adds, which could be attributed to the current implementation of the framework (TensorFlow) and can be improved by further parallel optimization.

We note that many base networks are stacked mainly with the permutation of 1×1 and 3×3 kernels and they can be applicable to embed AFDC in terms of the training and inference speed, i.e. [7, 11, 10, 32, 30] in ResNet stream and [9, 21, 34] in MobileNet stream. Besides, the adaptation is easy because our method is parameter-free. Overall, the random-size warping preserves the composition of the original image and also provides data augmentation to train the network with scale invariance. AFDC can adaptively construct fractional dilated kernels according to the spatial distortion information in a computation-efficient manner.

3.3. Composition-Aware Structure and Loss

The commonly-used network structures for the task of image classification usually incorporate global pooling before the fully connected layers [30, 7, 11, 26, 10]. The global pooling eliminates spatial variance which is helpful for the task of image recognition by training the networks with spatial invariant ability, but it causes the loss of localization information for image aesthetics assessment. Motivated by spatial pyramid pooling [6], MNA-CNN-Scene [16], several efforts are made to learn the information of spatial image compositions. First, we use multiple adaptive pooling modules [6] to output $g_i * g_i$ grids and feed them into the fully-connected layers (c.f. Fig. 2). The localization factors for image aesthetics assessment are highly correlated with the image symmetry and the overall image structure. Then, we aggregate the outputs after the fully-connected layers by concatenation. To limit the number of model parameters and prevent from overfitting, the module of each adaptive pooling layer outputs $\frac{num_features}{num_grids}$ channels.

Following the work in [27], we train our network to predict 10-scale score distribution with a softmax function on the top of the network. To get both the mean score prediction and the binary classification prediction, we calculate the weighted sum of score distribution $\sum_{i=1}^{10} i \cdot p_i$. We use the ordered distribution distance, Earth Mover Distance [27], as our loss function:

$$EMD(p, \hat{p}) = (\frac{1}{N} \sum_{k=1}^N |CDF_p(k) - CDF_{\hat{p}}(k)|^r)^{1/r} \quad (10)$$

where $CDF_p(k)$ is the cumulative distribution function as

network	cls. acc.	MSE	EMD	SRCC	LCC
NIMA(VGG16)[27]	0.8060	-	0.052	0.592	0.610
NIMA(Inception-v2)[27]	0.8151	-	0.050	0.612	0.636
NIMA(ResNet50, our implementation)	0.8164	0.3169	0.0492	0.6166	0.6388
Vanilla Conv (ResNet50)	0.8172	0.3101	0.0481	0.6002	0.6234
AFDC (random-size cropping pretrain)	0.8145	0.3212	0.0520	0.6134	0.6354
AFDC (aspect-ratio-preserving pretrain)	0.8295	0.2743	0.0445	0.6410	0.6653
AFDC + SPP	0.8324	0.2706	0.0447	0.6489	0.6711

Table 2: **Test result comparison on AVA [18]: The evaluation metrics are following [27]. Reported accuracy values(cls. acc.) are based on binary image classification. MSE(mean squared error), LCC (linear correlation coefficient) and SRCC (Spearman's rank correlation coefficient) are computed between predicted and ground truth mean scores. EMD measures the closeness of the predicted and ground truth rating distributions with $r = 1$ in Eq. (10). AFDC (random-size cropping) transfers the model trained with widely used data augmentation method in ImageNet, while AFDC (aspect-ratio-preserving pretrain) transfers the model trained with aspect-ratio-preserving data augmentation.**

$\sum_{i=1}^k p_i$. As stated in Section 1 and the results in [27], predicting the score distribution can provide more information about image aesthetics compared to the mean scores or binary classification labels.

4. Experimental Results

Following [27, 17, 15, 16, 12], we have evaluated our proposed method over AVA dataset [18]. The AVA contains around 250,000 images and each image contains the 10-scale score distribution rated by roughly 200 people. For a fair comparison, we use the same random split strategy in [27, 22, 17, 15, 16, 18] to generate 235,528 images for training and 20,000 images for test.

4.1. Implementation Details

We use ResNet-50 [7] as the backbone network due to its efficiency on computation and graphic memory as discussed in Section 3.2. We replace all the 3×3 Conv layers in each BottleNet with our proposed adaptive fraction dilation Conv layers. It is easy to plug AFDC into the common CNN architectures since it does not introduce any extra model parameters. We use the same EMD loss in Eq. (10) with $r = 2$ for better back propagation. To accelerate training, we use the grouping strategy discussed in Section 3.2. For the first 12 epochs, we train the model with three dilation kernels, 1×2 , 1×1 , 2×1 on the grouped images since the aspect ratios for 97.8% training and validation images fall between $[\frac{1}{2}, 2]$. Then we train the model with seven dilation kernels, 1×4 , 1×3 , 1×2 , 1×1 , 2×1 , 3×1 , 4×1 , for the remaining 6 epochs and select the best model from the results in the validation dataset. We note that the training and test speed could be further accelerated by a more fine-grained grouping strategy. We transfer the network parameters (pre-trained on ImageNet) before the fully connected layer and set the initial learning rate to 0.01 for the first 6 epochs. Then we dampen the learning rate to 0.001 for the rest of the training epochs. We find that

setting initial learning rate to 0.001 with a decay rate 0.95 after every 10 epochs can produce comparable results but converges more slowly. The weight and bias momentums are set to 0.9.

4.2. Ablation Study

In this section, we introduce the steps to build the final model and analyze the effects of each module step by step: (1) Replacing random cropping with composition-preserving random warping; (2) Replacing vanilla Conv with AFDC in the aspect-ratio-preserving pre-trained model on ImageNet; (3) Adding SPP modules to learn image composition.

Random Warping. For the data augmentation, input images in NIMA [27] are rescaled to 256×256 , and then a crop of size 224×224 is randomly extracted. They also report that training with random crops without rescaling produces the results that are not compelling due to the inevitable changes in image compositions. In order to preserve the complete composition, we replace the random-cropping with random-size warping by randomly warping each batch into square size in [224, 320] during each iteration. The network suffers from overfitting without using random warping. We note that non-square-size warping may further help with generalization and potentially train AFDC more robustly.

From Table 2, we generate slightly better results (Vanilla Conv (ResNet50)) compared with NIMA [27]. We use the same loss (EMD loss) and network (ResNet50, our implementation) as NIMA [27]. Comparable results have shown that random warping is an effective data augmentation alternative and it preserves the image composition.

Aspect-Ratio-Preserving Pretrain. We replace the vanilla convolution layers with AFDC in ResNet50. In our experiments, we find that, fine-tuning the fractional dilated convolution network results in similar validation accuracy compared to the original network (*c.f.* AFDC (random-size cropping pretrain) in Table 2). Compatible validation re-




Figure 6: The cropping results for the model trained with global pooling (left) and SPP (right). The two cropping samples are obtained by using a sliding window with the lowest score (green) and the highest score (red). The image is firstly resized to 256. A sliding window search with size 224 and stride 10 is applied.




Figure 7: The comparison of learning curves: the backbone networks here are all ResNet-50 [7].

sults might be attributed to the pre-trained model which has a distortion-invariant ability. The widely used data augmentation [26] for network training on ImageNet contains random cropping on a window whose size is distributed evenly between 8% to 100% of the original image area with the aspect ratio constrained to $[\frac{3}{4}, \frac{4}{3}]$. The model is trained with distortion invariance, which has the opposite interest of our method that tries to preserve the original aspect ratio.

For better transfer learning, we pre-train the ResNet50 [7] on ImageNet [4] without distortion augmentation. Specifically, we sample the 8% to 100% crop size to the image area with a square window, which is slightly modified comparing to the data augmentation method in [26]. As in Table 2, transferring the model from the aspect-ratio-preserving pre-train, we improve the overall test results (AFDC (aspect-ratio-preserving pre-train)) by a margin from the vanilla Conv counterpart.

Composition-Aware Structure. For better representation learning of composition, we use three different scales for SPP, $\{1 \times 1, 2 \times 2, 3 \times 3\}$. The network with a global pooling layer is equivalent to using only one scale, 1×1 . From Table 2, the network with SPP modules (AFDC+SPP) generates better results comparing to the network with the global pooling layer (AFDC). The experimental results have shown that incorporating the localization information could benefit the learning of image compositions. In Fig. 6, the automatic cropping example demonstrates that the ability of

localization/composition discrimination is important to find a good cropping result when the global cue in each cropping box has a similar distribution (color, lighting *et al.*). The model leaned with SPP modules can infer cropping respecting the image compositions, *e.g.* the relative position of eye and face in the example. We also tried $num_{grids} = 5$ and found that the results were not compelling due to the overfitting from extra model parameters. Three different scales are quite consistent with the common aesthetic rules (global information, symmetrical composition in horizontal and vertical direction, the rules of the thirds).

4.3. Effectiveness of AFDC

Learning Representation and Generalization From the experiments in Fig. 7, we argue that preserving aspect ratio information is essential for learning photo aesthetics since our method not only improves the validation results but also improves the training results. Without extra learning parameters, AFDC improves both learning representation and generalization ability. As discussed in Section 1, preserving the image aesthetics information completely omits the label noises caused by random warping and thus facilitates the learning process. The additional aesthetic features related to the aspect ratios allow the model to be more robust and discriminative. To further probe the effects of embedding aspect ratio, we compare different ways to incorporate the dilated convolution and the results are reported in Table 3. When trained with vanilla Conv (top rows in Table 3), AFDC is superior to other dilated Conv methods during the test. It implies the potential optimal between nearest two integer dilated kernels. After training with AFDC (bottom rows in Table 3), it further validates the effectiveness of AFDC, which is guided by the helpful supervision of aspect ratios. We note that such experiments are accessible because our method is parameter-free.

Overall, our proposed AFDC can learn more discriminative and accurate representations related to aesthetics perception, resulting in better generalization by leveraging extra supervision from the information of image aspect ratios.

Discriminative to Aspect Ratios To further investigate the response to aspect ratios, we resize the same image into different aspect ratios and test the results on different trained models. As shown in Fig. 8, AFDC (blue line) is discriminative to the change of aspect ratios. The small fluctuation of vanilla Conv (green line) is attributed to sampling change from resizing process. The model with random-size cropping pretrain on Imagenet (orange line) is less discriminative to capture the aesthetics perception related to aspect ratio due to its distortion-invariant pretrain. Moreover, the proposed method produces a multi-modal score distribution, which reflects that it learns complex relation between the aspect ratio and the aesthetics perception. It is in line with the notion that designing better aspect ratios or finding

Train	Test	cls.acc.	MSE	EMD
vanilla	vanilla	0.8172	0.3101	0.0481
	constant dilation rate = [2,1]	0.8072	0.5163	0.0610
	second nearest integer dilation	0.8091	0.5368	0.0620
	mean of nearest two integer dilations	0.8117	0.4558	0.0576
	nearest integer dilation	0.8114	0.4322	0.0562
	adaptive fractional dilation	0.8132	0.4133	0.0553
AFDC	vanilla	0.8085	0.3210	0.0581
	constant dilation rate = [2,1]	0.8132	0.3182	0.0576
	second nearest integer dilation	0.8156	0.3003	0.0476
	mean of nearest two integer dilations	0.8274	0.2771	0.0457
	nearest integer dilation	0.8277	0.2757	0.0457
	adaptive fractional dilation	0.8295	0.2743	0.0445

Table 3: **The test result comparison of different convolutions: The results are obtained with trained parameters by vanilla Conv (above) and AFDC (below). Test processes are conducted by different calculation methods for interpolation weights, w in Eq. (5). Vanilla Conv, constant dilation, nearest integer dilation and second nearest integer dilation can be interpreted as feeding one-hot interpolation weight vector into the networks.**




Figure 8: **Comparison of discrimination to the change of aspect ratios.**

aesthetically pleasing photography angles is not trivial.

Due to the constraint of training dataset, we admit that the learned perception related to the aspect ratios is not satisfactory yet even the model learns from different aspect ratios. As a matter of factor, the learning ability is available for our proposed method when training on a more specific targeted dataset. It could be utilized in automatic/auxiliary photo enhancement with not only color space transformation but also with spatial transformation, *e.g.* profile editing, multi-shot selection and automatic resizing.

4.4. Comparison With the State-of-the-Art Results

We have compared our adaptive fractional dilated CNN with the state-of-the-art methods in Table 4. The results of these methods are directly obtained from the corresponding papers. As shown in Table 4, our proposed AFDC outperforms other methods in terms of cls.acc and MSE, which are the most widely targeted metrics. Compared with NIMA(Inception-v2) [27] which uses the same EMD loss, our experimental results have shown that preserving the image aesthetic information completely results in better performance on image aesthetics assessment. We follow the same motivation from MNA-CMM-Scene [16], while our

Method	cls. acc.	MSE	SRCC
MNA-CNN-Scene [16]	76.5%	-	-
Kong <i>et al.</i> [14]	77.3%	-	0.558
AMP [17]	80.3%	0.279	0.709
Zeng <i>et al.</i> (resnet101) [33]	80.8%	0.275	0.719
NIMA (Inception-v2) [27]	81.5%	-	0.612
MP-Net [15] (50 cropping patches)	81.7%	-	-
Hosu <i>et al.</i> [8] (20 cropping patches)	81.7%	-	0.756
A-Lamp [15] (50 cropping patches)	82.5%	-	-
MP_{ada} [23](≥ 32 cropping patches)	83.0%	-	-
ours (single warping patch)	82.98%	0.273	0.648
ours (4 warping patches)	83.24%	0.271	0.649

Table 4: **Comparison with the SOTA methods: The four patches are warping size $\{224, 256, 288, 320\}$. The single patch is warping size 320 selected from the best results.**

proposed method is applicable to mini-batch training which contains images with different aspect ratios. The experimental results have shown adaptive embedding at kernel level is an effective way to learn more accurate aesthetics perception. Compared with multi-patch based methods [15, 8, 23], our unified model, which learns the image aesthetic features directly from the complete images in an end-to-end manner, can better preserve the original aesthetic information and alleviate the efforts to aggregate sampling prediction, *e.g.* complicated path sampling strategy and manually designed aggregation structure in [15]. Moreover, our method is much more efficient without feeding multiple cropping patches sampled from original images and could be more applicable for the application. Furthermore, it is much succinct due to its parameter-free manner and can be easily adapted to popular CNN architectures.

5. Conclusion

In this paper, an adaptive dilated convolution network is developed to explicitly model aspect ratios for image aesthetics assessment. Our proposed method does not introduce extra model parameters and can be plugged into popular CNN architectures. Besides, a grouping strategy has been introduced to reduce computational overhead. Our experimental results have demonstrated the effectiveness of our proposed approach. Even our adaptive dilated convolution network was proposed to support image aesthetics assessment, it can also be applied in other scenarios when image cropping or warping may introduce label noises. Moreover, adaptive kernel construction in a parameter-free manner provides an intuitive approach to design dynamic embedding at kernel level, which aims at better learning representation and generalization.

Acknowledgment

We would like to thank the anonymous reviewers for their helpful comments. This work was supported in part by NSFC under Grant (No. 61906143 and No.61473091).

References

- [1] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2018.
- [2] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017.
- [3] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Ze Wang. Studying aesthetics in photographic images using a computational approach. In *Computer Vision - ECCV 2006, 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006, Proceedings, Part III*, pages 288–301, 2006.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. Ieee, 2009.
- [5] Sagnik Dhar, Vicente Ordonez, and Tamara L. Berg. High level describable attributes for predicting aesthetics and interestingness. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*, pages 1657–1664, 2011.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *European conference on computer vision*, pages 346–361. Springer, 2014.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [8] Vlad Hosu, Bastian Goldlücke, and Dietmar Saupe. Effective aesthetics prediction with multi-level spatially pooled features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9375–9383, 2019.
- [9] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017.
- [10] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [11] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, volume 1, page 3, 2017.
- [12] Yueying Kao, Kaiqi Huang, and Steve Maybank. Hierarchical aesthetic quality assessment using deep convolutional neural networks. *Signal Processing: Image Communication*, 47:500–510, 2016.
- [13] Yan Ke, Xiaoou Tang, and Feng Jing. The design of high-level features for photo quality assessment. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006), 17-22 June 2006, New York, NY, USA*, pages 419–426, 2006.
- [14] Shu Kong, Xiaohui Shen, Zhe L. Lin, Radomír Mech, and Charles C. Fowlkes. Photo aesthetics ranking network with attributes and content adaptation. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I*, pages 662–679, 2016.
- [15] Shuang Ma, Jing Liu, and Chang Wen Chen. A-lamp: Adaptive layout-aware multi-patch deep convolutional neural network for photo aesthetic assessment. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 722–731, 2017.
- [16] Long Mai, Hailin Jin, and Feng Liu. Composition-preserving deep photo aesthetics assessment. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 497–506, 2016.
- [17] Naila Murray and Albert Gordo. A deep architecture for unified aesthetic prediction. *CoRR*, abs/1708.04890, 2017.
- [18] Naila Murray, Luca Marchesotti, and Florent Perronnin. Ava: A large-scale database for aesthetic visual analysis. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2408–2415. IEEE, 2012.
- [19] Masashi Nishiyama, Takahiro Okabe, Imari Sato, and Yoichi Sato. Aesthetic quality classification of photographs based on color harmony. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*, pages 33–40, 2011.
- [20] Jian Ren, Xiaohui Shen, Zhe L. Lin, Radomír Mech, and David J. Foran. Personalized image aesthetics. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 638–647, 2017.
- [21] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018* [10], pages 4510–4520.
- [22] Katharina Schwarz, Patrick Wieschollek, and Hendrik P. A. Lensch. Will people like your image? learning the aesthetic space. In *2018 IEEE Winter Conference on Applications of Computer Vision, WACV 2018, Lake Tahoe, NV, USA, March 12-15, 2018*, pages 2048–2057, 2018.
- [23] Kekai Sheng, Weiming Dong, Chongyang Ma, Xing Mei, Feiyue Huang, and Bao-Gang Hu. Attention-based multi-patch aggregation for image aesthetic assessment. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 879–886, 2018.
- [24] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [25] Xiaoshuai Sun, Hongxun Yao, Rongrong Ji, and Shaohui Liu. Photo assessment based on computational visual atten-

- tion model. In *Proceedings of the 17th International Conference on Multimedia 2009, Vancouver, British Columbia, Canada, October 19-24, 2009*, pages 541–544, 2009.
- [26] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [27] Hossein Talebi and Peyman Milanfar. NIMA: neural image assessment. *IEEE Trans. Image Processing*, 27(8):3998–4011, 2018.
- [28] Hanghang Tong, Mingjing Li, HongJiang Zhang, Jingrui He, and Changshui Zhang. Classification of digital photos taken by photographers or home users. In *Advances in Multimedia Information Processing - PCM 2004, 5th Pacific Rim Conference on Multimedia, Tokyo, Japan, November 30 - December 3, 2004, Proceedings, Part I*, pages 198–205, 2004.
- [29] Weining Wang, Mingquan Zhao, Li Wang, Jiexiong Huang, Chengjia Cai, and Xiangmin Xu. A multi-scene deep learning model for image aesthetic evaluation. *Sig. Proc.: Image Comm.*, 47:511–518, 2016.
- [30] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5987–5995. IEEE, 2017.
- [31] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [32] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference 2016, BMVC 2016, York, UK, September 19-22, 2016*, 2016.
- [33] Hui Zeng, Zisheng Cao, Lei Zhang, and Alan C Bovik. A unified probabilistic formulation of image aesthetic assessment. *IEEE Transactions on Image Processing*, 29:1548–1561, 2019.
- [34] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018* [10], pages 6848–6856.