

Revisiting Image Aesthetic Assessment via Self-Supervised Feature Learning

Kekai Sheng,^{1,2} Weiming Dong,^{1*} Menglei Chai,³ Guohui Wang,³ Peng Zhou,⁴
Feiyue Huang,² Bao-Gang Hu,¹ Rongrong Ji,⁵ Chongyang Ma⁶

¹Institute of Automation, Chinese Academy of Sciences, ²Youtu Lab, Tencent, ³Snap Inc.

⁴North China Electric Power University, ⁵Xiamen University, China, ⁶Kuaishou Technology
{saulsheng, garyhuang}@tencent.com, {weiming.dong, baogang.hu}@ia.ac.cn, {cmlatsim, robertwgh}@gmail.com
pzhou@ncepu.edu.cn, rrji@xmu.edu.cn, chongyangma@kwai.com

Abstract

Visual aesthetic assessment has been an active research field for decades. Although latest methods have achieved promising performance on benchmark datasets, they typically rely on a large number of manual annotations including both aesthetic labels and related image attributes. In this paper, we revisit the problem of image aesthetic assessment from the self-supervised feature learning perspective. Our motivation is that a suitable feature representation for image aesthetic assessment should be able to distinguish different expert-designed image manipulations, which have close relationships with negative aesthetic effects. To this end, we design two novel pretext tasks to identify the types and parameters of editing operations applied to synthetic instances. The features from our pretext tasks are then adapted for a one-layer linear classifier to evaluate the performance in terms of binary aesthetic classification. We conduct extensive quantitative experiments on three benchmark datasets and demonstrate that our approach can faithfully extract aesthetics-aware features and outperform alternative pretext schemes. Moreover, we achieve comparable results to state-of-the-art supervised methods that use 10 million labels from ImageNet.

Introduction

With the explosive growth of online visual data, the demand for image aesthetic assessment (Datta et al. 2006) in many multimedia applications has been dramatically increased. Typically, this assessment process seeks to evaluate the aesthetic level of each image according to certain rules commonly agreed by human visual perception, ranging from fine-grained local textures and lighting details to high-level semantic layout and composition. These highly subjective and ambiguous perceptual metrics pose formidable challenges on designing intelligent agents to automatically and quantitatively measure image aesthetics, especially with conventional hand-crafted features.

Following the recent advances in deep convolutional neural networks, researchers have explored various data-driven learning based approaches for aesthetic assessment and have reported impressive results in the past few years (Ma, Liu,




Figure 1: Illustration of our key idea. Some image editing operations typically have negative aesthetic impact and the degree of the impact is related to manipulation parameters.

and Chen 2017; Mai, Jin, and Liu 2016; Sheng et al. 2018b; Talebi and Milanfar 2018), benefiting from several image aesthetic benchmarks (Murray, Marchesotti, and Perronnin 2012; Kong et al. 2016). However, the inherent shortcomings of these datasets are still deterring us from continuously scaling up the volume of training data and improving the performance: (1) the sizes of existing image aesthetic datasets are far from enough to feed up latest neural networks with very deep architectures, since the labor work involved in manual labeling is prohibitively expensive and developing aesthetics-invariant data augmentation solution remains an open problem; (2) meanwhile, subjective human annotations are often strongly biased towards personal aesthetic preference, and thus require excessive amount of data to neutralize such inconsistency for reliable training.

Due to the poor scalability and consistency of available aesthetics datasets, it is becoming more and more attractive to break these bottlenecks of purely supervised training with unsupervised (Hinton and Salakhutdinov 2006; Vincent et al. 2008), or more specifically, self-supervised features (de Sa 1994; Doersch, Gupta, and Efros 2015; Pathak et al. 2016; Zhang, Isola, and Efros 2016). The main idea of self-supervision is to design pretext tasks that are naturally available, contextually relevant, and capable of providing proxy loss as the training signal to guide the

*Corresponding author

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

learning of features which are consistent with the real target. As a result, we can significantly broaden the scope of training data without introducing any further annotation cost.

In the context of image aesthetic assessment, due to the absence of an in-depth understanding of the human perception system, designing such a proper self-supervised pretext task can be challenging. However, although it is difficult to quantitatively measure the aesthetic score, predicting the relative influence of certain controlled image editing operations can be much easier. Therefore, we propose two key observations to help design the task (see Fig. 1): (1) some parametric image manipulation operations, such as blurring, pixelation, and patch shuffling, are very likely to have consistent negative impact on image aesthetics; (2) and the degree of impact due to such degradation will be monotonically increasing with respect to values of the corresponding operation parameters.

In this work, motivated by the strong correlation between image aesthetic levels and some degradation operations, we propose, to the best of our knowledge, the very first self-supervised learning scheme for image aesthetic assessment. The core idea behind our approach is to extract aesthetics-aware features with two novel self-supervision pretext tasks on distinguishing the type and strength of image degradation operations. To improve the training efficiency, we also introduce an entropy-based weighting strategy to filter out image patches with less useful training signals. The experimental results demonstrate that our self-supervised aesthetics-aware feature learning is able to achieve promising performance on available aesthetics benchmarks such as AVA (Murray, Marchesotti, and Perronnin 2012) and AADB (Kong et al. 2016), and outperforms a wide range of commonly adopted self-supervision schemes, including context prediction (Doersch, Gupta, and Efros 2015), jigsaw puzzle (Noroozi and Favaro 2016), colorization (Larsson, Maire, and Shakhnarovich 2017; Zhang, Isola, and Efros 2016), and rotation recognition (Gidaris, Singh, and Komodakis 2018).

In summary, our main contributions include:

- We propose a simple yet effective self-supervised learning scheme to extract useful features for image aesthetic assessment without using manual annotations.
- We present an entropy-based weighting strategy to help strengthen meaningful training signals in manipulated image patches.
- On three image aesthetic assessment benchmarks, our approach outperforms other self-supervised counterparts and even works better than models pre-trained on ImageNet or Places datasets using a large number of labels.

Related Work

Image aesthetic assessment has been extensively studied in the past two decades. Conventional solutions use hand-crafted feature extractors designed with domain expertise to model the aesthetic aspects of images (Datta et al. 2006; Ke, Tang, and Jing 2006). More recently, the power of deep neural networks makes it possible to learn feature representations that can surpass hand-crafted ones. Typi-

cal approaches include distribution-based objective functions (Talebi and Milanfar 2018; Jin et al. 2018), multi-level spatially pooling (Hosu, Goldlucke, and Saupe 2019), attention-based learning schemes (Sheng et al. 2018b), and attribute/semantics-aware models (Ma, Liu, and Chen 2017; Pan, Wang, and Jiang 2019). The advance of learning-based image aesthetic assessment has also inspired a number of practical solutions for various usage scenarios such as clothing (Yu et al. 2018) and food (Sheng et al. 2018a).

Self-supervised feature learning can be considered as one type of unsupervised learning algorithms (de Sa 1994), which intends to learn useful representations without manual annotations. Many effective pretext tasks have been proposed in this direction, such as context prediction (Doersch, Gupta, and Efros 2015), colorization (Zhang, Isola, and Efros 2016), split-brain (Zhang, Isola, and Efros 2017), and RotNet (Gidaris, Singh, and Komodakis 2018). The representations learned from self-supervised schemes turn out to be useful for many downstream tasks, *e.g.*, tracking (Vondrick et al. 2018), re-identification (Fan et al. 2018), and image generation (Lucic et al. 2019).

Aesthetics-aware image manipulations can be generally divided into two categories: (1) Rule-based approaches leverage empirical knowledge and domain expertise to enforce well-established photographic heuristics such as gamma correction and histogram equalization. (2) Data-driven approaches focus on learning powerful feature representations from examples to improve aesthetics in certain aspects, such as aesthetics-aware image enhancement (Deng, Loy, and Tang 2018; Hu et al. 2018), image blending (Hung et al. 2018), and colorization (Zhang, Isola, and Efros 2016).

Our Method

Key Observations

Our self-supervised feature learning approach for image aesthetic assessment is based on two key observations. First, experiments in previous work on image aesthetic assessment benchmarks indicate that inappropriate data augmentation (*e.g.*, brightness/contrast/saturation adjustment, PCA jittering) during the training process will result in performance degradation at the test time (Murray, Marchesotti, and Perronnin 2012; Kong et al. 2016). Second, it is observed that a convolutional neural network (CNN) model trained from manually annotated aesthetic labels can inherently acquire the ability to distinguish fine-grained aesthetic differences caused by various image manipulation methods. In Fig. 2, for instance, some image editing operations (such as Gaussian blur, downsampling, color quantization, and rotation) can increase the prediction confidence of aesthetically negative images, or even turn an aesthetically positive example into a negative one (as shown in the last row of Fig. 2).

Consequently, we argue that these fine-grained perceptual quality problems, without manual annotations, are closely related to the image aesthetic assessment task, for which meaningful training instances can be constructed via proper expert-designed image manipulations.

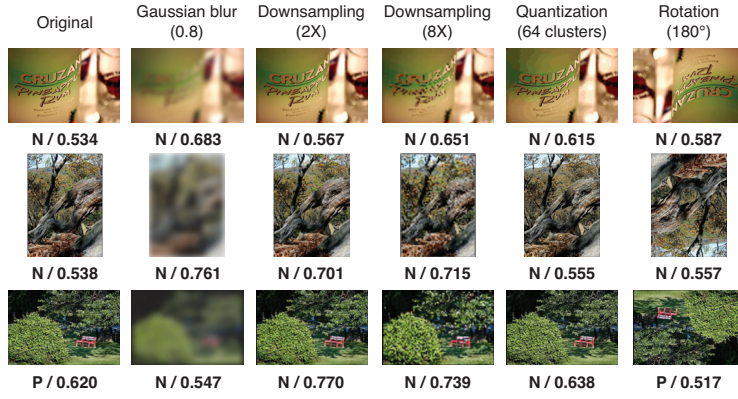


Figure 2: Some learning-free parametric image editing operations can introduce controllable aesthetic degradations on input images. The aesthetic label (**P** for positive and **N** for negative), as well as the corresponding assessment confidence predicted by a fully supervised model are listed beneath each image.

Table 1: Image editing operations and the parameters adopted to construct meaningful pretext tasks.

Attribute	Operation	Parameters
Much noise	JPEG compression	{60, 10}
	Gaussian noise	{0.2, 0.8}
Camera Shake	Rotation	{90, 180, 270}
Soft / Grainy	Downsampling	{4, 6}
	Quantization	{64, 8}
	Pixelation	{4, 8}
Poor lighting	Exposure	{0.5, 3.0}
Fuzzy	Gaussian blur	{0.2, 0.8}
Distracting	Patch shuffle	{0.1, 0.5}
	Mixup	{0.1, 0.4}

Selected Image Manipulations

According to empirical knowledge, different image editing operations have diverse effects on the manipulated output (Ma et al. 2018; Zhang et al. 2018b). Furthermore, some operations require complex parameter settings (e.g., inpainting) and it is often difficult to automatically compare the output to the input in terms of aesthetic level (e.g., grayscale conversion). In our case, however, we need to select image manipulations with easily controllable parameters and predictable perceptual quality.

Specifically, as listed in Table 1, we adopt a variety of image manipulation operations with different parameters for artificial training instances, including (1) downsampling by a scaling factor and upsampling back to the original resolution via bilinear interpolation; (2) JPEG compression with a percentage number to control the quality level; (3) Gaussian noise controlled by the variance; (4) Gaussian blur controlled by the standard deviation; (5) color quantization into a small number of levels; (6) brightness change based on a scaling factor; (7) random patch shuffle (Kang et al. 2017); (8) pixelation based on a patch size; (9) rotation by a certain degree; and (10) linear blending (mixup (Zhang et al. 2018a)) based on a constant alpha value.

Aesthetics-Aware Pretext Tasks

In this section, we propose our aesthetics-aware pretext task in a self-supervised learning scheme from two aspects. On one hand, the ability to categorize different types of image manipulations can be beneficial to the learning of aesthetics-aware representation. On the other hand, for the same type of editing operation, different control parameters can render various aesthetic levels correspondingly, and the tendency of quality shift is predictable (Marchesotti, Murray, and Perronnin 2015). Take JPEG compression for instance, decreasing the output image quality parameter always decreases the aesthetic level. Therefore, by constructing images with manipulations resulting in predictable degradation behaviors, we can extract meaningful training signals for the fine-grained aesthetics-aware pretext tasks.

Degradation identification loss. We denote an image patch as p , and the manipulation parameter as θ_t for $m(\cdot, \theta_t)$. The loss term of our first pretext task, *i.e.*, $\mathcal{L}_{\text{deg}}(p, t)$, reinforces the model to recognize which operation t has been applied to p :

$$\begin{aligned}\mathcal{L}_{\text{deg}}(p, \theta_t) &= -\log P_t(p; \mathbf{W}), \\ P_t(p; \mathbf{W}) &= P(\tilde{t} = t \mid m(p, \theta_t); \mathbf{W}),\end{aligned}\quad (1)$$

where $m(p, \theta_t)$ is the transformed output patch given the image patch p by the parameters θ_t , and $P(\tilde{t} = t \mid m(p, \theta_t); \mathbf{W})$ is the probability predicted by our model \mathbf{W} that p has undergone a degradation operation of type \tilde{t} .

For a comprehensive coverage of image attributes (e.g., resolution, color, spatial structures), we leverage a variety of typical manipulation operations T as listed in Table 1. To take better advantage of synthesized instances, we adopt parameters that will always induce aesthetic degradation of observably different patterns. Apart from these distortions, *None* operation is also taken into consideration. That is, we require the model to categorize 22 different classes of editing operations.

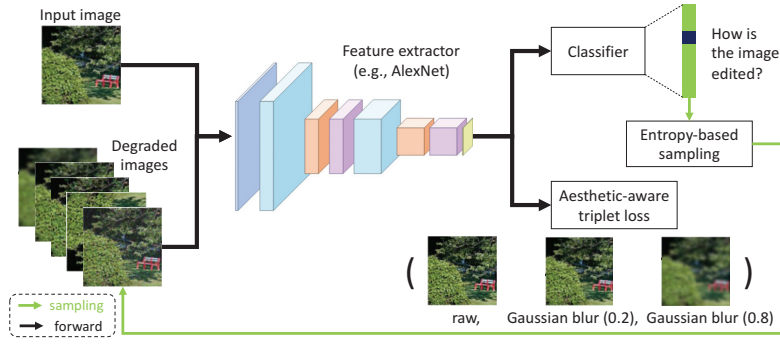


Figure 3: The diagram of our self-supervised approach. We propose two aesthetic-aware pretext tasks, and apply an entropy-based weighting scheme to enhance the training efficiency. In this way, we learn useful aesthetic-aware features in a label-free manner.

Triplet loss. Training with $\mathcal{L}_{\text{deg}}(p, t)$ alone is not enough for the task of aesthetic assessment, since some editing operations may produce some low-level artifacts that are easy to detect and will fool the network to learn some trivial features of no practical semantics (Gidaris, Singh, and Komodakis 2018). Our solution to address this issue is to encode the information via triplets $(p, m(p, \theta_{t_1}), m(p, \theta_{t_2}))$, where t_1 and t_2 are two different parameters of a certain operation in Tab. 1 (except for rotation and exposure). The two parameters are specified to create aesthetic distortions with a predictable relationship, *i.e.*, the edited image patch $m(p, \theta_{t_1})$ using t_1 is aesthetically more positive than $m(p, \theta_{t_2})$. Therefore, in an ideal aesthetic-aware representation, the distance between the original image patch p and $m(p, \theta_{t_1})$ should be smaller than the distance between p and $m(p, \theta_{t_2})$. In this way, we propose the second task, $\mathcal{L}_{\text{trp}}(p, \theta_{t_1}, \theta_{t_2})$:

$$\mathcal{L}_{\text{trp}}(p, \theta_{t_1}, \theta_{t_2}) = \max\{0, \mathcal{D}(p, \theta_{t_1}) - \mathcal{D}(p, \theta_{t_2}) + 1\}, \quad (2)$$

$$\mathcal{D}(p, \theta_t) = \left\| h(p, \mathbf{W}) - h(m(p, \theta_t), \mathbf{W}) \right\|_2^2,$$

where $h(p, \mathbf{W})$ is the normalized feature extracted from the model \mathbf{W} given a patch p .

It should be noted that we do not apply the triplet loss term from the beginning of the training process, since the representation learned from $\mathcal{L}_{\text{deg}}(p, t)$ in the early stage can oscillate drastically and thus may not be suitable for comparisons of fine-grained features. To combat the training dynamics, we activate $\mathcal{L}_{\text{trp}}(p, \theta_{t_1}, \theta_{t_2})$ after the curve of $\mathcal{L}_{\text{deg}}(p, t)$ has shown some plateau.

Total loss function. By putting the degradation identification loss item and the triplet loss item together, we formulate a new self-supervised pretext task as below:

$$\mathcal{L} = \frac{1}{|B| \cdot |T|} \sum_{p \in B} \sum_{t, t_1, t_2 \in T} \mathcal{L}_{\text{deg}}(p, \theta_t) + \lambda \cdot \mathcal{L}_{\text{trp}}(p, \theta_{t_1}, \theta_{t_2}), \quad (3)$$

where B is a mini-batch made up of $|B|$ patches, T is the set of all the image manipulations, and λ is a scaling factor to balance the two terms. The diagram of our proposed self-supervised scheme is shown in Fig. 3.

Entropy-based weighting. To improve the training efficiency and reduce noisy signals from misleading instances, we present a simple yet effective entropy-based weighting strategy. Specifically, after warm up for several epochs, we apply an entropy-based weight w_p for each patch p :

$$w_p = \max\{1 + \alpha \cdot \sum_t P_t(p; \mathbf{W}) \cdot \log P_t(p; \mathbf{W}), 0\},$$

where α is used to control the lower bound of w_p and we set it as $5/3$ in our experiments. The rationale behind this strategy is that instances of higher entropy values tend to have more uncertain visual cues for image aesthetics, and thus should be assigned with lower weights in the optimization process.

Discussions. Our proposed pretext task is similar to (Dosovitskiy et al. 2014; Liu, van de Weijer, and Bagdanov 2017) in which CNNs are trained to discriminate instances from different types of data augmentation methods. Different from these previous methods, we select a set of image manipulation operations and design our loss function carefully so that the learned representation can be aware of significant patterns of visual aesthetics. Besides, instead of building on selected high quality photos, we conduct pretext task optimization directly on images from ImageNet (Deng et al. 2009). This strategy makes our learning scheme more flexible to use.

One might argue that we ignore some global aesthetic factors, *e.g.*, rule of thirds. The reasons are two folds. First, global image attributes are more complex to manipulate than local ones which are considered in our approach. Second, global factors generally involve semantics which are not available in the context of self-supervised feature learning. We are not intended to mimic the statistics of *real-world* visual aesthetics, but to propose pretext tasks which are suitable for visual aesthetic assessment.

Experiments

Baseline Methods

We compare the performance of our method with five typical self-supervised visual pretext tasks, as listed below:

Context: The context predictor (Doersch, Gupta, and Efros 2015) that predicts the relative positions between two square patches cropped from one input image.

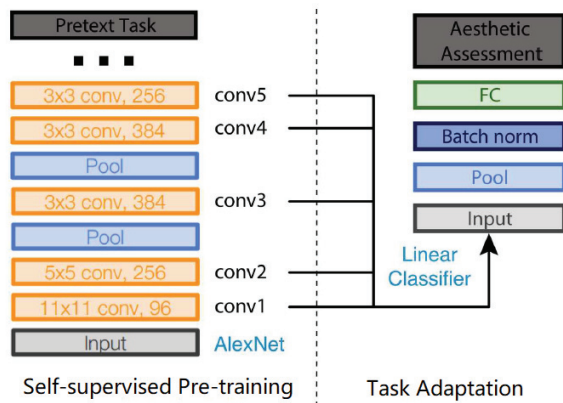


Figure 4: The schematic of a self-supervised learning framework, which we apply in our experiments to evaluate various pretext tasks for image aesthetic assessment.

Colorization: The image colorization task (Zhang, Isola, and Efros 2016) requires the model to estimate the color channels from a gray-scale image.

Split-brain: The cross-channel predictor (Zhang, Isola, and Efros 2017) estimates one subset of the color channels from another with constrained spatial regions.

Counting: The primitive counting task (Noroozi, Pirsiavash, and Favaro 2017) requires the model to generate visual primitives with their number invariant to transformations including both scaling and tiling.

RotNet: A rotation predictor (Gidaris, Singh, and Komodakis 2018) is trained to recognize the 2D rotation applied to input images, i.e., $m(\cdot; y)$ is image rotation and $y \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$.

In our experiments, we use the pre-trained models released by the authors for reliable and fair comparisons¹²³⁴⁵.

Additionally, we also testify with three typical pre-training strategies, including the method pre-trained with 1,000-way object labels from ImageNet (Deng et al. 2009) or 365-way scene labels from Places (Zhou et al. 2017), and the Gaussian random initialization i.e., without any pretext task.

Training Pipeline

Our entire training pipeline (Fig. 4) contains two parts: a *self-supervised pre-training* stage to learn visual features with unlabeled images and a *task adaptation* stage to evaluate how the learned features perform in the task of image aesthetic assessment.

In the pre-training stage (Fig. 4 left), the first few layers share the same structure with AlexNet (Krizhevsky, Sutskever, and Hinton 2012) for fair comparisons. In the

task adaptation stage (Fig. 4 right), following the same configurations in (Zhang, Isola, and Efros 2017), we freeze each model learned in the pre-training and leverage similar linear classifiers to evaluate visual features from each convolutional layer for the task of binary aesthetic classification. The channel number of each convolutional layer is shown in Fig. 4, and the dimensions of the corresponding fully connected layers are 9600×2 , 9216×2 , 9600×2 , 9600×2 , and 9216×2 , respectively.

Implementation details. In the pre-training stage, we first resize the shorter edge of each input image to 256. Then we randomly crop one patch of resolution 227×227 from the resized image. Next, we randomly choose three manipulation operations in Tab. 1 to edit each patch. We apply SGD optimization using a batch size of 64, with the Nesterov momentum of 0.9 and the weight decay of $5e-4$. We begin with a learning rate of 0.1, dropped it by a factor of 0.2 after every 10 epochs. To eschew training oscillating, we activate $\mathcal{L}_{trp}(p, t)$ with λ of 0.02 after the first 30 epochs. The following adaptation stage shares the same settings except that the learning-rate starts from 0.01.

Benchmarks for Aesthetic Assessment

Aesthetic Visual Analysis (AVA). The AVA dataset (Murray, Marchesotti, and Perronnin 2012) contains approximately 250,000 images. Each image has about 100 crowd-source aesthetic ratings in the range of 1 to 10. Following the common practice in (Ma, Liu, and Chen 2017; Hosu, Goldlucke, and Saupe 2019), we consider images whose average aesthetic scores are no less than 5.0 as positive instances and adopt the same training/test partition, i.e., 230,000 images for training and 20,000 for testing.

Aesthetics with Attributes Database (AADB). The AADB dataset (Kong et al. 2016) contains 10,000 images with aesthetic ratings $\{r \in [0, 1]\}$ and eleven additional attributes. We follow (Kong et al. 2016) to split the dataset into three partitions, i.e., 8500, 500, and 1000 images for training, validation, and testing, respectively. Without loss of generality, we binarize the aesthetic ratings into two classes using a threshold of 0.5, similar to AVA.

Chinese University of Hong Kong-Photo Quality Dataset (CUHK-PQ). The CUHK-PQ dataset (Luo, Wang, and Tang 2011) contains 17,690 images with binary aesthetic labels. Commonly used training/testing partitions on this dataset include a random 50/50 split and a five-fold split for cross-validation. We use the former one in our experiments.

Results and Discussions

Evaluation of Unsupervised Features

Our experimental results on the three benchmarks are reported in Table 2. In each column, the best numbers are shown in **bold** font and the second best are highlighted with an underscore. We can make several interesting observations from the table.

¹Context: <https://github.com/cdoersch/deepcontext>.

²Colorization: <https://github.com/richzhang/colorization>.

³Split-brain: <https://github.com/richzhang/splitbrainauto>.

⁴Counting: <https://github.com/gitlimlab/Representation-Learning-by-Learning-to-Count>.

⁵RotNet: <https://github.com/gidariss/FeatureLearningRotNet>.

Table 2: Task generalization performance (%) of different convolutional layers from models guided by different pretext tasks, measured for visual aesthetic assessment with linear layers on the AVA, AADB, and CUHK-PQ benchmarks.

Method	AVA						AADB						CUHK-PQ					
	conv1	conv2	conv3	conv4	conv5	average	conv1	conv2	conv3	conv4	conv5	average	conv1	conv2	conv3	conv4	conv5	average
ImageNet label	79.3	79.0	79.3	79.1	79.4	79.22	63.6	65.0	64.2	67.5	64.2	64.9	77.8	83.3	84.1	84.2	83.1	82.5
Places label	79.6	79.5	79.2	79.6	79.8	79.54	61.8	63.2	63.4	65.0	65.2	63.72	76.0	82.2	82.5	83.1	81.6	81.08
Without pretext	77.2	77.8	78.0	78.2	78.0	77.84	61.3	58.9	62.1	61.6	64.0	61.58	73.6	72.0	73.3	74.6	73.1	73.32
Context	79.8	79.2	79.2	79.0	79.0	79.24	59.4	59.8	62.4	62.6	63.2	61.48	69.2	75.9	77.9	79.0	78.7	76.16
Colorization	80.0	79.7	79.5	79.2	79.2	79.52	57.0	63.2	66.4	63.4	63.4	62.68	73.5	79.3	80.6	82.4	82.6	79.68
Split-Brain	79.5	79.7	79.5	80.1	79.4	79.64	58.2	64.4	67.8	63.8	65.6	63.96	77.5	83.7	84.7	83.7	84.5	82.82
Counting	53.0	52.2	63.3	65.5	58.8	58.56	<u>61.8</u>	61.3	60.4	62.8	62.3	61.72	75.6	76.0	74.3	72.5	71.8	74.04
RotNet	77.6	73.8	80.3	80.3	80.3	78.46	54.6	57.8	64.2	<u>66.0</u>	66.0	61.72	67.2	66.8	63.6	63.6	63.6	64.96
Ours	78.4	79.9	<u>80.5</u>	80.8	80.6	80.02	62.6	65.9	68.4	68.9	<u>65.8</u>	66.32	<u>77.4</u>	<u>83.4</u>	85.3	85.6	85.1	83.36




Figure 5: The accuracy results of the 3rd *conv.* block from different learning schemes in low data adaptation on AADB.

The proposed scheme generally works the best on all the three benchmarks. It is evident that our approach can achieve competitive results consistently, compared with other baselines on the three datasets, especially for the mid-level layers, *e.g.*, *conv4*. Self-supervised visual features can even outperform semantics-related features that are pre-trained with manual labels from ImageNet or Places. Furthermore, from the accuracy perspective, our method can be comparable to or work better than existing self-supervised learning schemes in image aesthetic assessment.

Mid-level features achieve the best results in task adaptation. By comparing performance of different layers, we can see the correlation between image aesthetics and network depth. As shown in all tables, mid-level features (*conv3* & *conv4*) generally outperform high-level features (*conv5*) and low-level ones (*conv1* & *conv2*) in terms of accuracy. The observations are consistent with (Zhang, Isola, and Efros 2016; Gidaris, Singh, and Komodakis 2018). One possible explanation is that, during the back propagation process, *conv5* gets more training signals for the pretext task as compared to *conv4*. Consequently, *conv5* is more likely to suffer from the overfitting issue, while using *conv4* leads to better generalization performance.

In addition to image attributes used in our pretext tasks, color is another important factor in image aesthetic assessment. Among other tested pretext tasks, Split-brain (Zhang, Isola, and Efros 2017) and Colorization (Zhang, Isola, and Efros 2016) consistently achieve the better results. It indicates




Figure 6: The accuracy results of the 3rd *conv.* block from different methods in low data adaptation on CUHK-PQ.

that color and image attributes in Tab. 1 are the key factors in assessing visual aesthetics. Regarding why RotNet fails to achieve good results on the CUHK-PQ benchmark, we suspect that both high-quality images and low-quality ones in this dataset share similar distributions or visual patterns.

Low Data Adaptation

One practical issue that self-supervised learning schemes are able to handle is low data adaptation, where very few labels are available for task adaptation. In our case, we simulate the low data adaptation regime by using 5% ~ 50% of the original training data per class in the task adaptation stage. The final results are shown in Fig. 5 and Fig. 6. From these two figures, we can see that our proposed learning scheme generally outperforms baseline models pre-trained with 1,000-ways labels from ImageNet. As the margin becomes larger when using more manual aesthetic labels in the task adaptation, our method presents higher efficiency of data usage compared to the vanilla supervised counterpart.

It is also interesting to note that our method and vanilla supervised method have similar adaptation performance when 1% ~ 10% training data is available. We believe this fact is due to the complexity of visual aesthetics, *i.e.*, when aesthetic labels are extremely few, the sampled instances cannot cover the entire distribution faithfully and thus lead to poor assessment results.

Table 3: Results of several methods on AVA benchmark measured in terms of binary classification accuracy of aesthetic labels.

Method	Backbone	Labels in the pre-training	Aesthetic label alone	Results (%)
DMA-Net (Lu et al. 2015)	AlexNet		×	75.4
AA-Net (Wang and Shen 2017)	VGG		×	76.9
(Kong et al. 2016)	AlexNet		×	77.3
MNA-CNN (Mai, Jin, and Liu 2016)	VGG	$\sim 10M$	×	77.4
NIMA (Talebi and Milanfar 2018)	Inception		✓	81.5
Pool-3FC (Hosu, Goldlucke, and Saupe 2019)	Inception		✓	81.7
A-Lamp (Ma, Liu, and Chen 2017)	VGG		×	82.5
MP_{ada} (Sheng et al. 2018b)	ResNet-18		✓	83.0
Our pretext task	AlexNet	0	✓	82.0
+ non-linear layers	ResNet-18			82.8

Adaption Using a Non-Linear Classifier

If we use non-linear layers in the task adaptation stage, we can achieve results close to state-of-the-art fully supervised approaches (Ma, Liu, and Chen 2017; Sheng et al. 2018b; Talebi and Milanfar 2018; Hosu, Goldlucke, and Saupe 2019) which have a test accuracy of 80% \sim 83% on the AVA benchmark, as shown in Tab. 3. Note that our method does not use 10 millions labels from ImageNet during the pre-training, and we only use aesthetic labels in the task adaptation stage. Besides, our accuracy is about 69.2% on the AADB benchmark and 88.5% on the CUHKPQ dataset, using ResNet-18 as the backbone network. These numbers are also close to that of the same network pre-trained with 10 millions labels from ImageNet (*i.e.*, *i.e.*, 70.4% on the AADB and 90.2% on the CUHKPQ). Arguably, we manage to achieve similar assessment performance using far less manual labels.

Ablation Study

Pretext tasks. We perform the pre-training with either one of the two pretext tasks and measure the final assessment results. It turns out that joint training using both loss terms generally yields better aesthetic assessment accuracy than using $\mathcal{L}_{deg}(p, \theta_t)$ or $\mathcal{L}_{trp}(p, \theta_{t_1}, \theta_{t_2})$ alone. The accuracy difference is $\sim 0.3\%$ on the AVA dataset and $\sim 0.7\%$ on the AADB dataset. We also find that pre-training with $\mathcal{L}_{trp}(p, \theta_{t_1}, \theta_{t_2})$ alone is prone to undesirable training dynamics, while the two-stage learning scheme is more stable and consequently leads to better results.

Image editing operations. We randomly select one operation in Tab 1 and exclude it from the pre-training stage to analyze the impact on the final assessment performance. We can make several observations from the corresponding results shown in Tab. 4. First, editing operations which are related to softness, camera shake, and poor lighting are relatively more significant in learning aesthetic-aware features, compared to other operations. Second, image manipulations which are related to distracting, fuzziness, and noise have relatively smaller margins. This fact indicates that these operations may create aesthetically ambiguous instances that do not always provide consistent signals. Interestingly, camera shake has

Table 4: The degradation caused by discarding some image manipulation operations during the pretext task training.

Attribute	AVA	AADB
Much noise	$\sim 0.1\%$	$\sim 0.3\%$
Poor lighting	$\sim 0.3\%$	$\sim 0.4\%$
Soft / Grainy	$\sim 0.4\%$	0.5 \sim 0.8%
Fuzzy	$\sim 0.1\%$	$\sim 0.2\%$
Camera shake	0.3 \sim 0.5%	0.4 \sim 0.6%
Distracting	$\sim 0.2\%$	0.3 \sim 0.4%

the most significant impact on the AVA, while soft/grainy are the most important ones for the AADB.

Entropy-based weighting. Without our entropy-based weighting scheme, there will be at least 0.5% down in the performance after task adaptation. Besides, some undesired dynamics will occur during the training process. Therefore, we can strengthen meaningful training signals from manipulated instances and improve training efficiency by assigning an entropy-based weight to each patch.

Conclusions

In this paper, we propose a novel self-supervised learning scheme to investigate the possibility of learning useful aesthetic-aware features without manual annotations. Based on the correlation between negative aesthetic effects and several expert-designed image manipulations, we argue that an aesthetic-aware representation space should distinguish between results yielded by various operations. To this end, we propose two pretext tasks, one to recognize what kind of editing operation has been applied to an image patch, and the other to capture fine-grained aesthetic variations due to different manipulation parameters. Our experimental results on three benchmarks demonstrate that the proposed scheme can learn aesthetics-aware features effectively and generally outperforms existing self-supervised counterparts. Besides, we achieve results comparable to state-of-the-art supervised methods on the AVA dataset without using labels from ImageNet. We hope these findings can help us obtain a better understanding of image visual aesthetics and inspire future research in related areas.

Acknowledgements

This work was supported in part by National Natural Science Foundation of China under nos. 61832016, 61672520 and 61720106006, in part by Key R&D Program of Jiangxi Province (No. 20171ACH80022), in part by CASIA-Tencent Youtu joint research project.

References

- Datta, R.; Joshi, D.; Li, J.; and Wang, J. Z. 2006. Studying aesthetics in photographic images using a computational approach. In *ECCV*, 288–301.
- de Sa, V. R. 1994. Learning classification with unlabeled data. In *NIPS*, 112–119.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*, 248–255.
- Deng, Y.; Loy, C. C.; and Tang, X. 2018. Aesthetic-driven image enhancement by adversarial learning. In *ACM MM*, 870–878.
- Doersch, C.; Gupta, A.; and Efros, A. A. 2015. Unsupervised visual representation learning by context prediction. In *ICCV*, 1422–1430.
- Dosovitskiy, A.; Springenberg, J. T.; Riedmiller, M.; and Brox, T. 2014. Discriminative unsupervised feature learning with convolutional neural networks. In *NIPS*, 766–774.
- Fan, H.; Zheng, L.; Yan, C.; and Yang, Y. 2018. Unsupervised person re-identification: Clustering and fine-tuning. *TOMM* 14(4):83.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised representation learning by predicting image rotations. In *ICLR*.
- Hinton, G. E., and Salakhutdinov, R. R. 2006. Reducing the dimensionality of data with neural networks. *Science* 313(5786):504–507.
- Hosu, V.; Goldlucke, B.; and Saupe, D. 2019. Effective aesthetics prediction with multi-level spatially pooled features. In *CVPR*, 9375–9383.
- Hu, Y.; He, H.; Xu, C.; Wang, B.; and Lin, S. 2018. Exposure: A white-box photo post-processing framework. *TOG* 37(2):26.
- Hung, W.-C.; Zhang, J.; Shen, X.; Lin, Z.; Lee, J.-Y.; and Yang, M.-H. 2018. Learning to blend photos. In *ECCV*, 70–86.
- Jin, X.; Wu, L.; Li, X.; Chen, S.; Peng, S.; Chi, J.; Ge, S.; Song, C.; and Zhao, G. 2018. Predicting aesthetic score distribution through cumulative jensen-shannon divergence. In *AAAI*, 77–84.
- Kang, G.; Dong, X.; Zheng, L.; and Yang, Y. 2017. PatchShuffle Regularization. *arXiv preprint arXiv:1707.07103*.
- Ke, Y.; Tang, X.; and Jing, F. 2006. The design of high-level features for photo quality assessment. In *CVPR*, 419–426.
- 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.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*, 1097–1105.
- Larsson, G.; Maire, M.; and Shakhnarovich, G. 2017. Colorization as a proxy task for visual understanding. In *CVPR*, 6874–6883.
- Liu, X.; van de Weijer, J.; and Bagdanov, A. D. 2017. RankIQ: Learning From Rankings for No-Reference Image Quality Assessment. In *ICCV*, 1040–1049.
- 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.
- Lucic, M.; Tschannen, M.; Ritter, M.; Zhai, X.; Bachem, O.; and Gelly, S. 2019. High-fidelity image generation with fewer labels. In *ICML*, 4183–4192.
- Luo, W.; Wang, X.; and Tang, X. 2011. Content-based photo quality assessment. In *ICCV*, 2206–2213.
- Ma, K.; Liu, W.; Zhang, K.; Duanmu, Z.; Wang, Z.; and Zuo, W. 2018. End-to-end blind image quality assessment using deep neural networks. *TIP* 27(3):1202–1213.
- Ma, S.; Liu, J.; and Chen, C. W. 2017. A-Lamp: Adaptive layout-aware multi-patch deep convolutional neural network for photo aesthetic assessment. In *CVPR*, 722–731.
- Mai, L.; Jin, H.; and Liu, F. 2016. Composition-preserving deep photo aesthetics assessment. In *CVPR*, 497–506.
- Marchesotti, L.; Murray, N.; and Perronnin, F. 2015. Discovering beautiful attributes for aesthetic image analysis. *IJCV* 113(3):246–266.
- Murray, N.; Marchesotti, L.; and Perronnin, F. 2012. AVA: A large-scale database for aesthetic visual analysis. In *CVPR*, 2408–2415.
- Noroozi, M., and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 69–84.
- Noroozi, M.; Pirsiavash, H.; and Favaro, P. 2017. Representation learning by learning to count. In *ICCV*, 5898–5906.
- Pan, B.; Wang, S.; and Jiang, Q. 2019. Image aesthetic assessment assisted by attributes through adversarial learning. In *AAAI*, 679–686.
- Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *CVPR*, 2536–2544.
- Sheng, K.; Dong, W.; Huang, H.; Ma, C.; and Hu, B.-G. 2018a. Gourmet photography dataset for aesthetic assessment of food images. In *SIGGRAPH Asia 2018 Technical Briefs*.
- Sheng, K.; Dong, W.; Ma, C.; Mei, X.; Huang, F.; and Hu, B.-G. 2018b. Attention-based multi-patch aggregation for image aesthetic assessment. In *ACM MM*, 879–886.
- Talebi, H., and Milanfar, P. 2018. NIMA: Neural image assessment. *TIP* 27(8):3998–4011.
- Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*, 1096–1103.
- Vondrick, C.; Shrivastava, A.; Fathi, A.; Guadarrama, S.; and Murphy, K. 2018. Tracking emerges by colorizing videos. In *ECCV*, 391–408.
- Wang, W., and Shen, J. 2017. Deep cropping via attention box prediction and aesthetics assessment. In *ICCV*, 2186–2194.
- Yu, W.; Zhang, H.; He, X.; Chen, X.; Xiong, L.; and Qin, Z. 2018. Aesthetic-based clothing recommendation. In *WWW*, 649–658.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018a. mixup: Beyond empirical risk minimization. In *ICLR*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018b. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 586–595.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *ECCV*, 649–666.
- Zhang, R.; Isola, P.; and Efros, A. A. 2017. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 1058–1067.
- Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Places: A 10 million image database for scene recognition. *TPAMI* 40(6):1452–1464.